Topological circles and Euler tours in locally finite graphs

Agelos Georgakopoulos*

Mathematisches Seminar
Universität Hamburg
Bundesstraße 55
20146 Hamburg, Germany

Submitted: Jan 15, 2008; Accepted: Mar 16, 2009; Published: Mar 25, 2009
Mathematics Subject Classification: 05C45

Abstract

We obtain three results concerning topological paths ands circles in the end compactification $|G|$ of a locally finite connected graph G. Confirming a conjecture of Diestel we show that through every edge set $E \in \mathcal{C}$ there is a topological Euler tour, a continuous map from the circle S^1 to the end compactification $|G|$ of G that traverses every edge in E exactly once and traverses no other edge.

Second, we show that for every sequence $(\tau_i)_{i\in\mathbb{N}}$ of topological x–y paths in $|G|$ there is a topological x–y path in $|G|$ all of whose edges lie eventually in every member of some fixed subsequence of (τ_i). It is pointed out that this simple fact has several applications some of which reach out of the realm of $|G|$.

Third, we show that every set of edges not containing a finite odd cut of G extends to an element of $\mathcal{C}$.

1 Introduction

Erdős et al. [17] characterised the infinite graphs G containing an *Euler tour*, that is, a two-way infinite walk traversing every edge of G precisely once. Although their result is best possible, it is not really satisfying: graphs with more than two ends cannot have such an Euler tour, so we cannot generalise to arbitrary infinite graphs the well-known theorems about Euler tours in finite ones. However, Diestel and Kühn [13] proposed a new concept of an Euler tour, called a *topological Euler tour*, that does allow such theorems to generalise to arbitrary infinite graphs, at least locally finite ones.

A *topological Euler tour* of a locally finite (multi-)graph G is a continuous map $\sigma : S^1 \to |G|$ traversing every edge of G precisely once. Here, $|G|$ is the Freudenthal

*Supported by a GIF grant.

compactification [18] of G, a topological space consisting of G and its ends, and S^1 is the unit circle in the real plane. (See Section 2 for precise definitions.) Unlike the approach of [17], this concept makes it possible to have a topological Euler tour in a graph that has more than two ends. Moreover, topological Euler tours found an important application in [23].

The following theorem can be considered as a generalisation of Euler's theorem that a finite connected multigraph has an Euler tour if and only if every vertex has even degree.

Theorem 1.1 ([13])**.** *Let G be a connected locally finite multigraph. Then, G admits a topological Euler tour if and only if every finite cut of G is even.*

The concept of a topological Euler tour and Theorem 1.1 are part of a far-reaching research project initiated by Diestel that uses topological concepts to generalise some classical parts of finite graph theory to infinite graphs; see [10] or [11, Chapter 8.5] for an introduction. In this context, a *circle* in $|G|$ is a homeomorphic image of S^1. The *(topological) cycle space* $\mathcal{C}(G)$ of G is a vector space over $\mathbb{Z}_2$ built from edge sets of circles in $|G|$; see Section 2. This concept has allowed the generalisation to infinite graphs of most of the well-known theorems about the cycle space of a finite graph, most of which fail for infinite graphs if the usual finitary notion of the cycle space is applied [2, 3, 4, 5, 13, 24].

A result in this context, which is used in the proof of Theorem 1.1, is the following.

Theorem 1.2 ([13])**.** *Let G be a connected locally finite multigraph. Then, $E(G) \in \mathcal{C}(G)$ if and only if G admits a topological Euler tour.*

If $H \subseteq G$ are finite connected multigraphs and $E(H) \in \mathcal{C}(G)$ then it follows easily from (the finite version of) Theorem 1.2 that H admits an Euler tour. Subgraphs of infinite multigraphs however are harder to handle, and one reason why this is the case is that, in general, they have a different end-topology compared to the original graph. It is thus not surprising that although Theorem 1.2 was one of the first results in this research project, it still remained open to prove

Theorem 1.3. *Let $H \subseteq G$ be locally finite multigraphs such that the closure $\overline{H}$ of H in $|G|$ is topologically connected. Then $E(H) \in \mathcal{C}(G)$ if and only if H admits a topological Euler tour in $|G|$.*

Here a *topological Euler tour of H in $|G|$* is a continuous map from S^1 to $|G|$ traversing each edge of H precisely once and meeting no other edge at an inner point. Theorem 1.3 was conjectured by Diestel [9], and it is the main result of this paper. We prove it in Section 4.

Our next result is a lemma whose main idea has already appeared implicitly several times (e.g. in [24] and in [23, Theorem 4]) and which could have many further applications. Given a locally finite graph G and two vertices x, y in G, this lemma facilitates the construction of a topological x-y path contained in a fixed subspace of $|G|$. Unlike finite paths, such a topological path cannot be constructed step by step attaching one edge at a time, because the structure of a topological path can be arbitrarily complicated: in general, a topological path consists of double rays whose ordering in the path can be of any countable ordinal type (see for example the circle in Figure 1, in which the double rays are arranged like the rational numbers). The main idea behind our lemma is to construct such a topological path as a limit of x-y paths in (finite) subgraphs of G. It states that any infinite sequence of topological paths between two fixed vertices x, y in $|G|$ has a subsequence whose limit contains a topological x-y path. More precisely, given a sequence $E_1, E_2, \ldots$ of sets, let us write

$$\liminf(E_n) := \bigcup_{i\in\mathbb{N}} \bigcap_{j>i} E_j$$

for the set of elements eventually in E_n. For a subspace X of $|G|$ we write $E(X)$ for the set of edges contained in X. Similarly, for a topological path τ we write $E(\tau)$ for the set of edges contained in the image of τ.

Lemma 1.4. *Let G be a locally finite graph, let $x, y \in V(G)$ and let $(\tau_i)_{i\in\mathbb{N}}$ be a sequence of topological x–y paths in $|G|$. Then, there is an infinite subsequence $(\tau_{a_i})_{i\in\mathbb{N}}$ of (τ_i) and a topological x–y path σ in $|G|$ such that $E(\sigma) \subseteq \liminf(E(\tau_{a_i}))$. Moreover, if no τ_i traverses any edge more than once then $E(G) \setminus E(\sigma) \subseteq \liminf(E(G) \setminus E(\tau_{a_i}))$, no edge is traversed by σ more than once, and for every finite subset F of $E(\sigma)$ there is an $n \in \mathbb{N}$ such that the linear ordering of F induced by σ coincides with that induced by τ_{a_i} for every $a_i > n$.*

Lemma 1.4 is applied in [20] (see also [21]) in order to prove that the Dirichlet problem has a unique proper solution in a locally finite electrical network of finite total resistance. In Section 5 we will use it to obtain an elementary proof of the following result, which is one of the most important tools in this research project:

Theorem 1.5 ([14]). *For any locally finite graph G, every element of $\mathcal{C}(G)$ is a disjoint union of edge-sets of circles.*

One of the applications of Theorem 1.5 is in the proof of Theorem 1.2. Theorem 1.5 has a long history; its first proof, in [14], applied to non-locally-finite graphs as well and was quite technical. A simpler proof for locally finite graphs inspired by [33] appears in [11, Theorem 8.5.8], but this proof still uses a non-trivial fact from topology and another result of Diestel and Kühn [15], that if G is locally finite and $X \subseteq |G|$ is closed and connected then X is path-connected[1]. Our proof is elementary and based only on Lemma 1.4.

Having seen Theorem 1.1 and Theorem 1.3 it is natural to ask for a generalisation to infinite multigraphs of the following result of Jaeger. For a multigraph G and $H \subseteq E(G)$, we say that H *extends* to an element of $\mathcal{C}(G)$ if there is $H' \in \mathcal{C}(G)$ with $H \subseteq H'$.

Theorem 1.6 ([26]). *If G is a finite multigraph and $H \subseteq E(G)$ then H extends to an element of $\mathcal{C}(G)$ if and only if H contains no cut of G of odd cardinality.*

[1] After the writing of this paper was completed, a simpler proof of this fact was published in [32], and a further proof of Theorem 1.5 appeared in [29].

In Section 6 we prove this assertion for locally finite G and discuss a possible application to hamiltonicity problems.

Finally, in Section 7 we discuss extensions of the above results to non-locally-finite graphs and other spaces.

2 Definitions

Unless otherwise stated, we will be using the terminology of [11] for graph theoretical concepts and that of [1] for topological ones. Let $G = (V, E)$ be a *locally finite* multigraph — i.e. every vertex has a finite degree — fixed throughout this section.

A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*. A *tail of* the ray R is a final subpath of R. Two rays R, L in G are *equivalent* if no finite set of vertices separates them. The corresponding equivalence classes of rays are the *ends* of G. We denote the set of ends of G by $\Omega = \Omega(G)$.

Let G bear the topology of a 1-complex[2]. To extend this topology to Ω, let us define for each end $\omega \in \Omega$ a basis of open neighbourhoods. Given any finite set $S \subset V$, let $C = C_G(S, \omega)$, or just $C(S, \omega)$ if G is fixed, denote the component of $G - S$ that contains some (and hence a tail of every) ray in ω, and let $\Omega(S, \omega)$ denote the set of all ends of G with a ray in C. As our basis of open neighbourhoods of ω we now take all sets of the form

$$C(S, \omega) \cup \Omega(S, \omega) \cup E'(S, \omega) \tag{1}$$

where S ranges over the finite subsets of V and $E'(S, \omega)$ is any union of half-edges $(z, y]$, one for every S–C edge $e = xy$ of G, with z an inner point of e. Let $|G|$ denote the topological space of $G \cup \Omega$ endowed with the topology generated by the open sets of the form (1) together with those of the 1-complex G.

It can be proved (see [12]) that in fact $|G|$ is the Freudenthal compactification [18] of the 1-complex G.

A *circle* in $|G|$ is the image of a homeomorphism from S^1, the unit circle in $\mathbb{R}^2$, to $|G|$. An *a*rc in $|G|$ is a homeomorphic image of the real interval $[0, 1]$ in $|G|$.

A subset D of E is a *circuit* if there is a circle C in $|G|$ such that $D = \{e \in E | e \subseteq C\}$. Call a family $(D_i)_{i \in I}$ of subsets of E *thin*, if no edge lies in D_i for infinitely many i. Let the *sum* $\Sigma_{i \in I} D_i$ of this family be the set of all edges that lie in D_i for an odd number of indices i, and let the *cycle space* $\mathcal{C}(G)$ of G be the set of all sums of (thin families of) circuits.

A *topological Euler tour of* G is a continuous map $\sigma : S^1 \to |G|$ such that every inner point of an edge of G is the image of exactly one point of S^1 (thus, every edge is traversed exactly once, and in a "straight" manner).

A *trail* in a multigraph G is a walk in G that traverses no edge more than once.

[2]Every edge is homeomorphic to the real interval $[0, 1]$, the basic open sets around an inner point being just the open intervals on the edge. The basic open neighbourhoods of a vertex x are the unions of half-open intervals $[x, z)$, one from every edge $[x, y]$ at x.

If H is a finite submultigraph of the locally finite multigraph G, then let H^* denote the multigraph obtained from G by contracting each component C of $G - H$ into a single vertex u_C, deleting loops but keeping multiple edges. If X is a topological path in $|G|$ then let $X \restriction H$ be the walk induced in H^* by X (induced by replacing each maximal subpath of X in a component C of $G - H$ by u_C). If $X \restriction H$ traverses some edge more than once then pick a trail with edges in $E(X \restriction H)$ between the endvertices of $X \restriction H$ and denote it by $X|H$; if $X \restriction H$ is a trail then let $X|H := X \restriction H$.

3 Basic facts

The following basic fact can be found in [25, p. 208].

Lemma 3.1. *The image of a topological path with endpoints x, y in a Hausdorff space X contains an arc in X between x and y.*

The well known orthogonality relation between elements of the cycle space and cuts has also been extended to locally finite graphs and provides a very useful tool [11, Theorem 8.5.8]:

Lemma 3.2. *Let $F \subseteq E(G)$. Then $F \in \mathcal{C}(G)$ if and only if F meets every finite cut in an even number of edges.*

The following lemma, proved e.g. in [11, Theorem 8.5.5], says that an arc is not allowed to jump a finite cut:

Lemma 3.3. *If $F = E(X, Y)$ is a finite cut of G (where $\{X, Y\}$ is a bipartition of $V(G)$) then any arc in $|G|$ with one endpoint in X and one in Y contains an edge from F.*

4 Euler tours

In this section we prove Theorem 1.3, but before we do so let us consider an example indicating why, contrary to the case when G is finite, this is harder to prove than Theorem 1.2. Let me first sketch the proof of Theorem 1.2 [13]. The harder implication is to show that if $E(G) \in \mathcal{C}(G)$ then G has a topological Euler tour. To prove this, Diestel & Kühn [13] first applied Theorem 1.5 to obtain a decomposition of $E(G)$ into a set $\mathcal{B}$ of pairwise edge-disjoint circuits. Then, they picked one of those circuits B_0 and a homeomorphism σ_0 from S^1 to the defining circle of B_0. They proceeded inductively in infinitely many steps, in each step n considering an element B of $\mathcal{B}$ such that B is incident with a vertex v visited by σ_{n-1} but the edges of B are not traversed by σ_{n-1}. Then, they modified σ_{n-1} locally in order to obtain a new continuous map σ_n that traverses B as well as all elements of $\mathcal{B}$ traversed by σ_{n-1}. This process yielded a limit map σ whose image contained every edge exactly once, and the remaining task was to show that this map was continuous.

Now consider the graph F of Figure 1, a well known object from [13, 10]. Formally, it is defined as follows. Let V be the set of finite 01 sequences, including the empty sequence

$\emptyset$. Define a tree on V by joining every $\ell \in V$ to its two one-digit extensions, the sequences $\ell 0$ and $\ell 1$. For every $\ell \in V$, add another edge e_ℓ between the vertices $\ell 01$ and $\ell 10$, and let D_ℓ denote the double ray consisting of e_ℓ and the two rays starting at e_ℓ whose vertices have the form $\ell 1000\ldots$ and $\ell 0111\ldots$. Finally, let L be the double ray whose vertices are $\emptyset$ and the all-zero and the all-one sequences, and let $\mathcal{D} := \{L\} \cup \{D_\ell \mid \ell \in V\}$ be the set of all those double-rays (drawn thick in Figure 1). It has been shown in [13] that all the elements of $\mathcal{D}$ and all the (continuum many) ends of F together form a circle C in $|F|$ (the double rays D_l are traversed by C in the order they appear in Figure 1 when moving from left to right; one can think of C as being the boundary of the outer face of F).

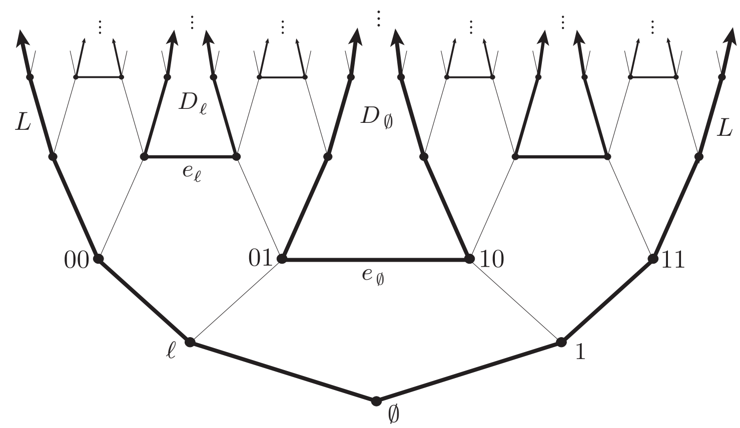

Figure 1: The graph F. The thick double-rays together with the ends of the graph form a circle.

The example we are going to use in order to see how the proof of Theorem 1.2 fails to prove Theorem 1.3 is the graph G obtained by taking two disjoint copies F_1, F_2 of F and joining every vertex of F_1 to its copy in F_2 by an edge. The subgraph H of $|G|$ we choose consists of the copies of $\mathcal{D}$ in both F_1, F_2. Its closure $\overline{H}$ contains all ends of G. Note that if R is a ray in F_1 and R' its copy in F_2, then R and R' belong to the same end of G. This implies that G and F have the same ends and $\Omega(G)$ and $\Omega(F)$ have the same topology; formally, the function $\pi : \Omega(F) \to \Omega(G)$ that maps every end of F to the end of G containing it as a subset is a homeomorphism. This means that the two copies C_1, C_2 of C in F_1, F_2 respectively are also circles in $|G|$, thus $\overline{H}$, which is the union of C_1 and C_2, is path-connected since C_1 and C_2 intersect: $C_1 \cap C_2 = \Omega(G)$.

Thus H satisfies the conditions of Theorem 1.3: we have $E(H) = E(C_1) \cup E(C_2) \in \mathcal{C}(G)$, and $\overline{H}$ is connected since it is path-connected. Now let us try to prove that H has a topological Euler tour in $|G|$ using the proof of Theorem 1.2 sketched above. Firstly, we have to apply Theorem 1.5. If we are lucky this will decompose $E(H)$ into $E(C_1)$ and $E(C_2)$, and we will be able to imitate that proof (although we will have to attach the two circles at an end rather than at a vertex). However, we can be unlucky: Theorem 1.5 could return a quite different decomposition into circuits. Note that for every $D \in \mathcal{D}$ the two copies of D in G together with their two incident ends form a circle C_D in $|G|$, comprising two double rays and two ends (recall that the two copies of a ray of F in G belong to the same end of G). Now $E(H)$ can be decomposed into the set of all edge-sets of such circles C_D for all $D \in \mathcal{D}$. But if the application of Theorem 1.5 returns this decomposition then we cannot continue with the proof: if D, D' are any two distinct elements of $\mathcal{D}$ then $\overline{D} \cap \overline{D'} = \emptyset$, so we cannot start with some circuit and then stepwise attach the others, since there are no common points where we could attach.

Thus we are going to need a different approach in order to prove our result:

Proof of Theorem 1.3. Let us start with the easier of the two implications: if H has a topological Euler tour in $|G|$ then, since arcs cannot jump finite cuts by Lemma 3.3, every finite cut of G has an even number of edges in H. Lemma 3.2 now implies that $E(H) \in \mathcal{C}(G)$.

We now assume conversely that $E(H) \in \mathcal{C}(G)$ and construct a topological Euler tour σ of H in $|G|$.

Let $v_1, v_2, \dots$ be an enumeration of the vertices of $V(G)$ and let G_n be the submultigraph $G[v_1, v_2, \dots, v_n]$ of G induced by the first n vertices in this enumeration. Moreover, let $H_n := (V(G_n^*), E(H) \cap E(G_n^*))$ be the submultigraph of G_n^* formed by the edges in $E(H) \cap E(G_n^*)$ and their incident vertices (the notation G_n^* is defined in Section 2).

Note that every cut of G_n^* is also a finite cut of G. Thus, if H_n intersects some cut of G_n^* in an odd number of edges, then so does H for the corresponding cut of G, which cannot be the case by Lemma 3.2. This means that for any bipartition $\{X, Y\}$ of $V(G_n^*)$, H_n contains an even number of X–Y edges. Moreover, H_n is connected, since otherwise there is a finite cut of G_n^*, and G, that separates H_n but contains no edge of H, and this cut easily contradicts the topological connectedness of $\overline{H}$. Thus by Euler's theorem [11] H_n has an Euler tour.

Note that every Euler tour W of H_{n+1} induces an Euler tour $W|n$ of H_n; indeed, let $W|n$ be the walk obtained from W by contracting all edges not in $E(H_n)$. By construction $W|n$ contains all edges of H_n and traverses each of them only once.

We want to apply König's Infinity Lemma [27] (or [11]) to obtain a sequence $(W_n)_{n\in\mathbb{N}}$ such that W_n is an Euler tour of H_n and $W_{n+1}|n = W_n$ for every n. So let V_n be the set of all Euler tours of H_n, and join any $W \in V_{n+1}$ to $W|n \in V_n$. The Infinity Lemma now yields a sequence $(W_n)_{n\in\mathbb{N}}$ with $W_n \in V_n$ as required (W_0 being the trivial walk).

We are now going to construct a topological Euler tour σ of H in $|G|$ as the limit of a sequence σ_n of topological Euler tours of the H_n corresponding to W_n.

Let σ_0 be the constant map from S^1 to the vertex u_C of G_0^* that corresponds to the component C of G that meets H. Assume inductively that for some n we have defined $\sigma_n : S^1 \to |G_n|$ so that σ_n runs through the edges and vertices of W_n in order, traversing every edge in a straight manner and pausing for a non-trivial interval at every vertex it visits (and each time it visits that vertex); more formally, the inverse image under σ_n of a vertex is always a disjoint union of non-trivial subarcs of S^1, one for each visit of W_n to that vertex.

We are now going to use σ_n in order to define σ_{n+1}. It is easy to see that W_n can be obtained from W_{n+1} by contracting every maximal subwalk that does not contain edges in W_n to a "contracted vertex"

$$u \in V(G_n^*) \backslash V(G_n). \tag{2}$$

For every such subwalk D, there is one subarc I_D of S^1 mapped to u by σ_n that corresponds to D: there is an occurrence of u in W_n that is split into two occurences of u in W_{n+1} between which D (and nothing else) is traversed, and I_D is the subarc that corresponds to that occurence.

Now modify σ_n so that it maps I_D continuously to D traversing every edge e of D precisely once, in a straight manner, and in the same direction as W_{n+1} traverses e, and pausing for a non-trivial interval at every vertex it visits. Let σ_{n+1} be the map obtained from σ_n by performing all these modifications for all contracted walks D (note that if D, D' are distinct maximal contracted walks then $I_D \cap I_{D'} = \emptyset$; thus all these modifications can be performed simultaneously).

We now define the limit map $\sigma : S^1 \to |G|$ as follows. Let $x \in S^1$ be given. If there is a point y so that $\sigma_n(x) = y$ for all but finitely many n, then put $\sigma(x) = y$. (In this case y must be a vertex of G or an inner point of an edge, and the maps σ_n will agree on x for all n so large that this vertex or edge is in G_n.) If not, then $\sigma_n(x)$ is a "contracted vertex" u_{C_n} for every n by (2), and the corresponding components C_n are nested. Let $\overline{C_n}$ denote the topological closure of C_n in $|G|$. As $|G|$ is compact, $U := \bigcap_n \overline{C_n}$ is non-empty, and as, clearly, $\bigcap_n C_n = \emptyset$, we have $U \subseteq \Omega(G)$. However, for any two distinct ends ω, ω' in $\Omega(G)$ there is a finite vertex set S separating their rays. This means that if $S \subseteq V(G_n)$ then $\overline{C_n}$ cannot contain both ω and ω', thus U consists of a single end ω. Let $\sigma(x) = \omega$. Note that

$$\sigma(x) \in \overline{C_n} \text{ for every } n. \tag{3}$$

This completes the definition of σ and now we have to show that it is continuous. This is trivial at points of S^1 that are mapped to some vertex or inner point of an edge, so let x be a point mapped to an end ω. We have to specify for every open neighbourhood O of ω an open subset of S^1 containing x and mapped to O by σ. However, for every basic open set $O' = O(S, \omega)$ of ω, there is an $n \in \mathbb{N}$ such that $\overline{C} \subseteq O'$ for some component C of $G - G_n$; just choose n large enough that $S \subseteq \{v_1, \dots, v_n\}$. This component C corresponds to a vertex $u_C \in V(G_n^*)$, and the inverse image of u_C under σ_n is, by the construction of σ_n, a disjoint union I of non-trivial subarcs of S^1. We claim that $\sigma(I) \subseteq \overline{C}$. Indeed, if $x \in I$ is mapped to a vertex or inner point of an edge by σ, then clearly $\sigma(x)$ lies in C. If x is mapped to an end, then $\sigma(x) \in \overline{C}$ holds by (3). This proves that σ is continuous. By construction σ traverses each edge in $E(H)$ precisely once and traverses no other edge, thus it is a topological Euler tour of H in $|G|$. □

5 An elementary proof of Theorem 1.5

In this section we prove Lemma 1.4 and apply it to obtain a proof of Theorem 1.5 which is shorter than previous proofs ([14, 11]).

Proof of Lemma 1.4. Let $v_1, v_2, \dots$ be an enumeration of the vertices of $V(G) - \{x, y\}$ and let G_n be the submultigraph $G[x, y, v_1, v_2, \dots, v_n]$ of G induced by x,y and the first n vertices in this enumeration. Given any topological path τ in $|G|$ let $\tau|n := \tau|G_n$. We begin by selecting the sequence (τ_{a_i}), in ω steps, choosing one member in each step. We will choose (τ_{a_i}) so that all members with index i or greater will agree on G_i^*:

$$\tau_{a_i}|i = \tau_{a_j}|i \text{ for all } i, j \in \mathbb{N} \text{ with } i < j. \tag{4}$$

For the first step, we pick an x–y trail W_1 in G_1^* such that the set $\mathcal{X}_1$ of indices j such that $\tau_j|1 = W_1$ is infinite; such a trail exists because G_1^* is finite and $\tau_i|1$ cannot use any edge more than once. Then, pick any element j of $\mathcal{X}_1$ and let $\tau_{a_1} := \tau_j$.

Having performed the first $n-1$ steps we pick, in step n, an x–y trail W_n in G_n^* such that the set $\mathcal{X}_n$ of indices $j \in \mathcal{X}_{n-1}$ such that $\tau_j|n = W_n$ is infinite; such a trail exists because G_n^* is finite and $\mathcal{X}_{n-1}$ infinite. Then, pick any element j of $\mathcal{X}_n$ and let $\tau_{a_n} := \tau_j$.

Thus we have defined the sequence $(\tau_{a_i})_{i\in\mathbb{N}}$, which by construction satisfies (4). We will now construct a topological path σ from x to y in $|G|$ with $E(\sigma) \subseteq \liminf(E(\tau_{a_i}))$ in ω steps. For the first step, let σ_0 be a continuous map from the real interval $[0,1]$ to G_1^* that runs through each edge of W_1 precisely once and in order, runs through no other edge, and maps a non-trivial interval I_v^j to each vertex v of G_1^* obtained by contracting a component of $G - G_1$ each time it visits v.

Then, at step n, modify σ_{n-1} on these intervals I_v^j to extend it to a continuous map σ_n from $[0,1]$ to G_n^* that runs through each edge of W_n precisely once and in order, also mapping non-trivial intervals I_v^j to each contracted vertex v it visits; it is possible to extend σ_{n-1} in this way by (4).

We can now define the required topological path $\sigma : [0,1] \to |G|$ as a "limit" of the σ_n: for every $p \in [0,1]$, if there is an $n_0 \in \mathbb{N}$ such that $\sigma_m(p) = \sigma_{n_0}(p)$ for every $m \geq n_0$, then we let $\sigma(p) = \sigma_{n_0}(p)$; otherwise, $\sigma_n(p)$ is a contracted vertex for every n, and as in the proof of Theorem 1.3 it is easy to prove that the components contracted to the $\sigma_n(p)$ converge to a unique end ω of G, and we let $\sigma(p) = \omega$.

Similarly with the proof of Theorem 1.3, it is straightforward to check that σ is continuous, thus it is a topological x–y path. By construction σ has the required properties. $\square$

Note that Lemma 1.4 remains true if one or both of x, y is an end. To see this, modify the proof of Lemma 1.4 so that W_n is an $\hat{x}$-$\hat{y}$ trail in G_n^*, where $\hat{x}$ is the contracted vertex of G_n^* corresponding to the component containing rays in x if x is an end, and $\hat{x} = x$ if x is a vertex; define $\hat{y}$ similarly.

Proof of Theorem 1.5. It suffices to prove that every non-empty element of $\mathcal{C}(G)$ contains at least one circuit; indeed, assume this is true, and for any $C \in \mathcal{C}(G)$ let Z be a maximal set of edge-disjoint circuits in C. Then, $C' = C + \sum_{D\in Z} D = C \setminus \bigcup_{D\in Z} D$ is an element of $\mathcal{C}(G)$, and if it is non-empty then it contains a circuit by our assumption, which contradicts the maximality of Z. Thus $C' = \emptyset$ and C is the sum of the family Z.

To show that any non-empty $C \in \mathcal{C}(G)$ contains a circuit, we will pick an arbitrary edge $f = xy \in C$ and find a circuit in C containing f. By definition, C is the sum of a thin family $\{E(C_i)\}_{i\in\mathbb{N}}$ of circuits of circles C_i. Let $F := \bigcup_{i\in\mathbb{N}} E(C_i)$, let $e_1, e_2, \ldots$ be an enumeration of the edges of G, and let $E_n := \{f\} \cup \{e_i \mid i < n, e_i \in F \setminus C\}$.

Lemma 5.1. *For every $n \in \mathbb{N}$ there is an x–y arc X_n in $|G|$ with $E(X) \subseteq F \setminus E_n$.*

Proof. Since $\{C_i\}$ is a thin family, the subset K_n of its elements that contain edges in E_n is finite. For every $C_i \in K_n$ replace every maximal subarc R of C_i that does not contain any edge in E_n by an edge e_R, called a *representing edge*, to obtain a finite cycle C'_i. Taking the union of all these new cycles we obtain a finite auxiliary graph

$$G_n := \Big(\bigcup_{C_i\in K_n} V(C'_i),\ \bigcup_{C_i\in K_n} E(C'_i)\Big)$$

and we can apply the finite version of Theorem 1.5 (see [11]) to the element $\sum_{C_i\in K_n} C'_i$ of $\mathcal{C}(G_n)$. This yields a disjoint union of finite circuits, and one of them, say D, has to contain f since f must appear in an odd number of elements of $\{C_i\}$ and thus of K_n. Moreover, D does not contain any edge in $E_n - f$, as those edges appear in an even number of elements of $\{C_i\}$ and thus of K_n. Now replacing every representing edge e_R in D by the arc R it represents and removing f yields a topological x–y path which, by Lemma 3.1, contains the required x–y arc. □

Applying Lemma 1.4 on the sequence $(\tau_n)_{n\in\mathbb{N}}$ where τ_n is a topological path whose image is the arc X_n obtained from Lemma 5.1 yields a topological x–y path all of whose edges lie in C, and applying Lemma 3.1 we obtain an x–y arc X. The union of X with f is the required circuit. □

6 Cyclability

In this section we generalise Theorem 1.6 to the topological cycle space of a locally finite graph:

Theorem 6.1. *If G is a locally finite multigraph and $H \subseteq E(G)$ then H extends to an element of $\mathcal{C}(G)$ if and only if H contains no cut of G of odd cardinality.*

In order to prove the backward implication of Theorem 6.1 we will first show that it is true for finite H, and then use this fact and compactness to prove it for any H.

For H, G as in Theorem 6.1 let $G[H]$ denote the submultigraph of G formed by the edges in H and their incident vertices.

Lemma 6.2. *For every locally finite multigraph G and every finite $H \subseteq E(G)$, if H contains no cut of G of odd cardinality then it extends to an element of $\mathcal{C}(G)$.*

Proof. We may assume without loss of generality that G is connected, since otherwise we can apply the result to each component of G that meets H.

We are going to find a finite submultigraph G' of G containing H such that H extends to an element of $\mathcal{C}(G')$. For this, consider every (possibly empty) $F \subseteq E(H)$, and if F

separates $G[H]$ but not G (in particular, if $F = \emptyset$ and $G[H]$ is not connected), then pick a set $\mathcal{P}_F$ of paths in G so that $(G[H] - F) \cup \bigcup \mathcal{P}_F$ is connected. Let G' be the (finite) union of $G[H]$ with all paths that lie in some $\mathcal{P}_F$. By the construction of G', if H contains a cut B of G' then B is also a cut of G. Thus H contains no cut of G' of odd cardinality, and we can apply Theorem 1.6 to H, G' to prove that H extends to an element of $\mathcal{C}(G')$. Since any cycle in G' is also a cycle in G, this means that H also extends to an element of $\mathcal{C}(G)$. □

We can now prove Theorem 6.1 applying Lemma 6.2 to the finite subsets of H and using compactness. One way to do this is using the methods of the previous sections and König's Infinity Lemma as the reader can check, however the most convenient way is to use the following compactness theorem for propositional logic (see [23] for a detailed explanation of our usage of Theorem 6.3):

Theorem 6.3. *Let K be an infinite set of propositional formulas every finite subset of which is satisfiable. Then K is satisfiable.*

Proof of Theorem 6.1. If H contains a cut of G of odd cardinality then by Lemma 3.2 it does not extend to an element of $\mathcal{C}(G)$.

For the reverse implication, let $e_1, e_2, \ldots$ be an enumeration of the elements of H, and for every $i \in \mathbb{N}$ let $H_i = H \cap \{e_1, \ldots, e_i\}$. By Lemma 6.2 there is for each i a $C_i \in \mathcal{C}(G)$ with $H_i \subseteq C_i$. In order to apply Theorem 6.3, we introduce for each edge e of G a boolean variable v_e; intuitively, the two possible values T, F that v_e can assume encode the presence or not of e. For every edge e in H let P_e denote the propositional formula $v_e = T$ (encoding the event that e is present). Moreover, for every finite cut B of G let P_B be a propositional formula (in the variables v_e) that evaluates to T if and only if the set $\{e \in B \mid v_e = T\}$ has even cardinality (P_B expresses the event that an even number of edges in B are present). Let $K := \{P_e \mid e \in H\} \cup \{P_B \mid B \text{ is a finite cut of } G\}$ be the set of all these propositional formulas.

We claim that every finite subset K' of K is satisfiable. Indeed, pick $n \in \mathbb{N}$ large enough that H_n contains all edges e such that $P_e \in K'$. Now C_n yields an assignment of values to the variables v_e that satisfies all formulas in K': for every $e \in E(G)$ let $v_e = T$ if $e \in C_n$ and $v_e = F$ otherwise. It is obvious that this assignment satisfies all formulas in K' of the form P_e. It also satisfies all formulas of the form P_B as, by Lemma 3.2, C_i meets every finite cut in an even number of edges.

We can thus apply Theorem 6.3 to obtain an assignment of values to the v_e that satisfies all formulas in K. The corresponding edge set $C := \{e \in E(G) \mid v_e = T\}$ contains, then, all edges in H, and it lies in $\mathcal{C}(G)$ by Lemma 3.2. □

Theorem 1.6 is proved by an elegant algebraic argument ([26]), and in fact this argument can be adapted to give an alternative proof to Theorem 6.1. Let G be a finite multigraph. For a subset F of $E(G)$, denote by $P(F)$ the subspace of the edge space $\mathcal{E}$ of G generated by the edges in F, and let $\bar{F} := E(G) \setminus F$. For $X \subseteq \mathcal{E}$ denote by $X^\perp$ the set of elements of $\mathcal{E}$ that meet each element of X in an even number of edges. Moreover, let K be the subspace of $\mathcal{E}$ consisting of the cuts of G, and let $\mathcal{C}$ be the cycle space of G. In this notation, Theorem 1.6 can be formulated as follows:

$$\text{If } G \text{ is a finite multigraph and } H \subseteq E(G) \text{ then } H \in \mathcal{C} + P(\bar{H}) \text{ if and only if } H \in [K \cap P(H)]^{\perp}. \tag{5}$$

Now to prove (5) note that $[\mathcal{C} + P(\bar{H})]^{\perp} = \mathcal{C}^{\perp} \cap P(\bar{H})^{\perp} = K \cap P(H)^{\perp}$, thus (5) follows from the fact that $X^{\perp\perp} = X$ (it is well-known that $\mathcal{C}^{\perp} = K$; see [11]). In order to apply the same argument in the case that G is an infinite locally finite multigraph, we need three facts (if X is a subspace of the edge space $\mathcal{E}$ of G then $X^{\perp}$ consists of those elements of $\mathcal{E}$ that have a (finite and) even intersection with each element of X). Firstly, we may assume without loss of generality that G is 2-edge-connected, and thus Lemma 3.2 implies that $\mathcal{C}^{\perp}$ is the space of finite cuts of G. Second, that $\mathcal{C}(G)$ is generated by a thin set (see again [11, Theorem 8.5.8]), and finally, that if $X \subseteq \mathcal{E}$ is generated by a thin set then $X^{\perp\perp} = X$, which follows from Theorem 4.3 of [7].

Let me now argue that Theorem 6.1 could find applications in the study of hamiltonicity in infinite graphs. If G is a locally finite graph then we define a *Hamilton circle* of G to be a circle in $|G|$ containing all vertices (and hence also all ends, as it is closed). This concept allowed the generalisation of Fleischner's theorem to locally finite graphs [23]. See [22] for more results and open problems about Hamilton circles.

Thomassen [31] conjectured that every finite 4-connected line graph is hamiltonian. An easy special case of this conjecture is to prove that the line graph of a 4-edge-connected finite graph G is hamiltonian. To prove this, note that since G is 4-edge-connected it has two edge-disjoint spanning trees T, S ([11, Theorem 2.4.1]). By Theorem 1.6, $E(T)$ extends to an element C of $\mathcal{C}(G)$: it contains no cut of G since its complement contains the spanning tree S. So let W be an Euler tour of $G[C]$, which exists by Theorem 1.2. Since W meets all vertices of G it is easy to convert W into a Hamilton cycle of the line graph $L(G)$ of G: we just translate edges of W into vertices of $L(G)$, and each time we visit a new vertex v of G while running along W we make a detour to pick up any edges incident with v that are not in W and have not been picked up yet; see [8] for a more precise exposition.

Having generalised Theorem 1.6 to Theorem 6.1 we could try to extend this proof to locally-finite graphs G. However the first step is still missing: we need a spanning subgraph of G that has the same end-topology as G and whose edge-set extends to an element of $\mathcal{C}(G)$ to play the role of T in the above sketch.

7 Non-locally-finite graphs and other spaces

In this section we are going to discuss generalisations of the results of this paper to non-locally-finite graphs and other spaces.

Let us start with an example. The graph G of Figure 2 consists of two edges $e = xu, f = vy$ and an infinite family $\{P_i\}_{i \in I}$ of indipendent u–v paths of length two. For every $i \in \mathbb{N}$ let R_i be the x–y path in G going through P_i. At first sight, this example seems to be suggesting that Lemma 1.4 cannot be extended to non-locally-finite graphs, since no u–v path appears in an infinite subsequence of $(R_i)_{i\in\mathbb{N}}$. However, we will see that obstructions of this kind can be overcome by choosing a topology on the graph in which pairs of points like u and v are *topologically indistinguishable*, that is, have precisely the same set of open neighbourhoods. The topology we are going to use, denoted by $||G||$, is well known (see [6, 21, 28, 30]) and resembles $|G|$, the main difference being that the role of finite vertex-separators in the definition of $|G|$ is now played by finite edge-separators. Formally, we define $||G||$, where G is an arbitrary graph, as follows.

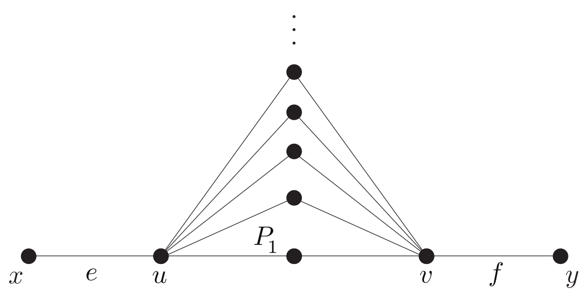

Figure 2: A simple non-locally-finite graph.

Call two rays in G *edge-equivalent* if no finite set of edges separates them, call the corresponding equivalence classes the *edge-ends* of G, and let $\Omega'(G)$ denote the set of edge-ends of G. (If G is locally finite then its edge-ends coincide with its ends.) We now endow the space consisting of G, considered as a 1-complex (i.e. every edge is a homeomorphic copy of the real interval $[0,1]$), and its edge-ends with the topology $||G||$. Firstly, every edge $e \in E(G)$ inherits the open sets corresponding to open sets of $[0,1]$. Moreover, for every finite edge-set $S \subset E(G)$, we declare all sets of the form

$$C(S,\omega) \cup \Omega'(S,\omega) \cup E'(S,\omega) \tag{6}$$

to be open, where $C(S,\omega)$ is any component of $G - S$ and $\Omega'(S,\omega)$ denotes the set of all edge-ends of G having a ray in $C(S,\omega)$ and $E'(S,\omega)$ is any union of half-edges $(z,y]$, one for every edge $e = xy$ in S with y lying in $C(S,\omega)$. Let $||G||$ denote the topological space of $G \cup \Omega'$ endowed with the topology generated by the above open sets.

Using the topology $||G||$ we just defined we can extend Lemma 1.4 to countable graphs:

Lemma 7.1. *Let G be a countable graph, let $x, y \in V(G) \cup \Omega'$ and let $(\tau_i)_{i\in\mathbb{N}}$ be a sequence of topological x–y paths in $||G||$. Then, there is an infinite subsequence $(\tau_{a_i})_{i\in\mathbb{N}}$ of (τ_i) and a topological x–y path σ in $||G||$ such that $E(\sigma) \subseteq \liminf(E(\tau_{a_i}))$. Moreover, if no τ_i traverses any edge more than once then $E(G) \setminus E(\sigma) \subseteq \liminf(E(G) \setminus E(\tau_{a_i}))$, no edge is traversed by σ more than once, and for every finite subset F of $E(\sigma)$ there is an $n \in \mathbb{N}$ such that the linear ordering of F induced by σ coincides with that induced by τ_{a_i} for every $a_i > n$.*

Proof (sketch). A proof can be obtained from the proof of Lemma 1.4 by slight modifications reflecting the fact that finite edge-separators now play the role that finite vertex-separators used to play. Here we only point out these modifications.

Firstly, the graphs G_i^* in that proof have to be defined differently now: we fix an enumeration $e_1, e_2, \dots$ of $E(G)$ and let G_i^* be the graph obtained from G by contracting all edges $e_{i+1}, e_{i+2}, \dots$. Note that if u, v are two vertices of G joined by an infinite set of pairwise edge-disjoint paths then u and v will be contracted into the same vertex in each of the G_i^*.

The graphs G_i^* are still finite, and each of the topological paths τ_i still induces a path in G_i^* like, in the proof of Theorem 1.4. The rest of the proof can be imitated, but some attention is required at the points $p \in [0,1]$ that are mapped by σ_i to a contracted vertex for every step i: it may be the case that there is no edge-end of G to which the images $(\sigma_i(p))_{i\in\mathbb{N}}$ converge, but in this case there is always a vertex v of infinite degree to which $(\sigma_i(p))_{i\in\mathbb{N}}$ converges, and we put $\sigma(p) = v$ for some such vertex v. If there are more than one candidate vertices for v then they are all topologically indistinguishable in $||G||$, so our choice will not matter for the continuity of σ. The proof that σ is continuous is still straightforward. $\square$

Following these lines and imitating the proofs of this paper the interested reader will be able to prove, with little effort, that Theorems 1.3 and 6.1 also extend to countable graphs using $||G||$ (Theorem 6.1 extends even to uncountable graphs). A generalisation to $||G||$ of Lemma 3.2, which is needed for the latter results, can be found in [30, Satz 4.3]. Theorem 1.5 has already been extended to arbitrary graphs, see [30, Satz 4.5].

The topological cycle space $\mathcal{C}$ of a locally finite graph bears some similarity to the first singular homology group H_1 of $|G|$, as both are based on circles, i.e. closed 1-simplices, in $|G|$. Diestel and Sprüssel [16] investigated the precise relationship between $\mathcal{C}(G)$ and $H_1(|G|)$, and found out that they are not the same. Still, a homology group was introduced in [19] that combines properties of singular homology and $\mathcal{C}$, and generalises Theorem 1.5 to arbitrary continua. The interested reader is referred to [19] or [21, Section 5].

References

[1] M.A. Armstrong. *Basic Topology.* Springer-Verlag, 1983.

[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)*, 92:235–256, 2004.

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

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

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

[6] D.I. Cartwright, P.M. Soardi, and W. Woess. Martin and end compactifications for nonlocally finite graphs. *Trans. Am. Math. Soc.*, 338:679–693, 1993.

[7] K. Casteels and B. Richter. The Bond and Cycle Spaces of an Infinite Graph. *J. Graph Theory*, 59(2):162–176, 2008.

[8] P. A. Catlin. Supereulerian graphs: A survey. *J. Graph Theory*, 16:177–196, 1992.

[9] R. Diestel. personal communication.

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

[11] R. Diestel. *Graph Theory* (3rd edition). Springer-Verlag, 2005. Electronic edition available at: `http://www.math.uni-hamburg.de/home/diestel/books/graph.theory`.

[12] R. Diestel and D. Kühn. Graph-theoretical versus topological ends of graphs. *J. Combin. Theory (Series B)*, 87:197–206, 2003.

[13] R. Diestel and D. Kühn. On infinite cycles I. *Combinatorica*, 24:68–89, 2004.

[14] R. Diestel and D. Kühn. On infinite cycles II. *Combinatorica*, 24:91–116, 2004.

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

[16] R. Diestel and P. Sprüssel. The homology of locally finite graphs with ends. preprint 2008.

[17] P. Erdős, T. Grünwald, and E. Vazsonyi. Über Euler-Linien unendlicher Graphen. *J. Math. Phys. Mass. Inst. Technology*, 17:59–75, 1938.

[18] H. Freudenthal. Über die Enden topologischer Räume und Gruppen. *Math. Zeitschr.*, 33:692–713, 1931.

[19] A. Georgakopoulos. Cycle decompositions in graphs and continua. In Preparation.

[20] A. Georgakopoulos. The Dirichlet problem in a network of finite total resistance. In Preparation.

[21] A. Georgakopoulos. Graph topologies induced by edge lengths. Preprint 2009.

[22] A. Georgakopoulos. Fleischner's theorem for infinite graphs. *Oberwolfach reports*, 2007.

[23] A. Georgakopoulos. Infinite hamilton cycles in squares of locally finite graphs. *Adv. Math.*, 220:670–705, 2009.

[24] A. Georgakopoulos and P. Sprüssel. Geodesic topological cycles in locally finite graphs. Submitted.

[25] D.W. Hall and G.L. Spencer. *Elementary Topology.* John Wiley, New York 1955.

[26] F. Jaeger. A note on subeulerian graphs. *J. Graph Theory*, 3:91–93, 1979.

[27] D. König. *Theorie der endlichen und unendlichen Graphen.* Akademische Verlagsgesellschaft, 1936.

[28] B. Krön. End compactifications in non-locally-finite graphs. *Math. Proc. Cambridge Phil. Soc.*, 131:427–443, 2001.

[29] R.B. Richter and A. Vella. Cycle Spaces in Topological Spaces. *J. Graph Theory*, 59(2):115–144, 2008.

[30] M. Schulz. Der Zyklenraum nicht lokal-endlicher Graphen (in german). Master's thesis, Universität Hamburg, 2005.

[31] C. Thomassen. Reflections on graph theory. *J. Graph Theory*, 10:309–324, 1986.

[32] C. Thomassen and A. Vella. Graph-like continua, augmenting arcs, and Menger's theorem. *Combinatorica*, 29. DOI: 10.1007/s00493-008-2342-9.

[33] A. Vella. *A fundamentally topological perspective on graph theory, PhD thesis*. Waterloo, 2004.