On Infinite Cycles II

To the memory of C.St.J.A. Nash-Williams

Reinhard Diestel Daniela Kühn

Abstract

We extend the basic theory concerning the cycle space of a finite graph to arbitrary infinite graphs, using as infinite cycles the homeomorphic images of the unit circle in the graph together with its ends. We characterize the spanning trees whose fundamental cycles generate the cycle space, and prove infinite analogues to the standard characterizations of finite cycle spaces in terms of edge-decomposition into single cycles and orthogonality to cuts.

1 Introduction

One of the basic and well-known facts about finite graphs is that their *fundamental* cycles C_e (those consisting of a chord $e = xy$ on some fixed spanning tree T together with the path xTy joining the endvertices of e in T) generate their entire cycle space: every cycle of the graph can be written as a sum mod 2 of fundamental cycles. Answering a question of Richter, we obtained in [5] a generalization of this fact to locally finite infinite graphs and a natural notion of infinite cycles. Our approach was to consider the compact topological space $\overline{G}$ of G together with its ends, and to define a *circle* in $\overline{G}$ as a homeomorphic image in $\overline{G}$ of the unit circle. Thus every finite cycle of G is a circle in $\overline{G}$, but $\overline{G}$ can also have infinite circles, i.e. circles containing infinitely many edges. It turns out that it makes sense to define a *cycle* of G to be the set of edges contained in a circle in $\overline{G}$, and hence we may extend the definition of the cycle space of a finite graph to infinite graphs in a way that allows for both infinite (topological) cycles and infinite sums (mod 2) of these cycles.

The main result of [5] is that for suitable spanning trees of a locally finite graph (namely, for its end-faithful spanning trees) all the cycles are generated by the (finite) fundamental cycles. The same result is true also for the other elements of the cycle space, those that are non-trivial sums of cycles. (This does not follow trivially for infinite sums.)

Our first aim in this paper is to prove similar results for arbitrary infinite graphs: we shall characterize the spanning trees whose fundamental cycles generate every cycle (Section 3) or the entire cycle space (Section 4).

Our second main aim will be to extend to infinite graphs – and infinite cycles – two further standard results about finite cycle spaces: that every element of the cycle space of a graph is an edge-disjoint union (rather than just a sum) of cycles, and that a subgraph is an element of the cycle space if and only if it meets every cut in an even number of edges. These two results will be proved in Section 5, which uses a couple of lemmas from Section 4 but can otherwise be can be read independently of Sections 3 and 4.

We remark that, for some graphs, it is possible to strengthen our results by allowing certain infinite sums in the definition of the cycle space that cannot be allowed in general (and will therefore be disallowed in this paper): sums where every edge lies in at most finitely many summands but some vertices may lie in infinitely many. Those extensions require adjustments to the end set of the underlying graph and its topology: only its 'topological ends' (see [7]) are added as new 'points at infinity', while rays from other ends converge to certain vertices. See [6] for details.

2 Basic facts and terminology

The terminology we use is that of [2], and we assume familiarity with [5]. We shall freely view a graph either as a combinatorial object or as the topological space of a 1-complex. So 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; note that we do not require local finiteness here.

A homeomorphic image (in the subspace topology) of the unit interval in a topological space X will be called an *arc* in X; a homeomorphic image of the unit circle in X is a *circle* in X. When A is an arc in X, we denote the set of all inner points of A by $\mathring{A}$. Similarly, when E is a set of edges, we write $\mathring{E}$ for the set of all inner points of edges in E. The following lemma can be proved by elementary topological arguments.

Lemma 2.1 *Every arc in G between two vertices is a graph-theoretical path. If X is an open subset of G, then the set of points in X that can be reached* *by an arc in X from a fixed point $x \in X$ is open. The topological components of X coincide with its arc-connected components.* $\square$

Given a spanning tree T in a graph G, every edge $e \in E(G) \setminus E(T)$ is a *chord* of T, and the unique cycle C_e in $T + e$ is a *fundamental cycle* with respect to T. A rooted spanning tree T of G is *normal* if the endvertices of every edge of G are comparable in the tree order induced by T. Countable connected graphs have normal spanning trees, but not all uncountable ones do; see [8] for details. We will use the following simple lemma, a proof can be found in [9].

Lemma 2.2 *Let $x_1, x_2 \in V(G)$, and let T be a normal spanning tree of G. For $i = 1, 2$ let P_i denote the path in T joining x_i to the root of T. Then $V(P_1) \cap V(P_2)$ separates x_1 from x_2 in G.* $\square$

We refer to 1-way infinite paths as *rays*, to 2-way infinite paths as *double rays*, and to the subrays of rays or double rays as their *tails*. If we consider two rays in a graph G as *equivalent* if no finite set of vertices separates them in G, then the equivalence classes of rays are known as the *ends* of G. (The grid, for example, has one end, and the binary tree has continuum many; see [3] for more background.) We shall write $\overline{G}$ for the union of G (viewed as a space, i.e. a set of points) and the set of its ends.

Given an end ω and a finite set S of vertices of G, there is exactly one component $C = C_G(S, \omega)$ of $G - S$ which contains a tail of every ray in ω. We say that ω *belongs* to C. Let $\overline{C}_G(S, \omega)$ denote the union of $C := C_G(S, \omega)$ with the set of all ends of G belonging to C. Write $E_G(S, \omega)$ for the set of all edges between S and C in G. Let Top denote the topology on $\overline{G}$ generated by the open sets of the 1-complex G and all sets of the form

$$\widehat{C}_G(S, \omega) := \overline{C}_G(S, \omega) \cup \mathring{E}'_G(S, \omega)\,,$$

where $\mathring{E}'_G(S, \omega)$ is any union of half-edges $(x, y] \subset e$, one for every $e \in E_G(S, \omega)$, with $x \in \mathring{e}$ and $y \in C$. So for each end ω, the sets $\widehat{C}_G(S, \omega)$ with S varying over the finite subsets of $V(G)$ are the basic open neighbourhoods of ω.

Throughout this paper we assume that $\overline{G}$ is endowed with Top. Thus $\overline{G}$ is a Hausdorff space in which every ray converges to the end that contains it. The following lemma summarizes some properties of arcs and circles in $\overline{G}$. The proof of the first part (for circles) can be found in [5, Lemma 4.3], the remainder can be proved by elementary (though not completely trivial) topological arguments.

Lemma 2.3 *For every arc A and every circle C in $\overline{G}$ the sets $A \cap G$ and $C \cap G$ are dense in A and C, respectively. Moreover, every arc A in $\overline{G}$ whose endpoints are vertices or ends, and every circle C in $\overline{G}$, includes every edge of G of which it contains an inner point. If x is a vertex in $\mathring{A}$ (respectively on C), then A (respectively C) contains precisely two edges of G at x.* □

By Lemma 2.3 every circle in $\overline{G}$ 'has' a well-defined set of edges, and it can be recovered from those edges as their closure in $\overline{G}$. It therefore makes sense to define a *cycle* in G as a subgraph consisting of all the edges contained in a given circle in $\overline{G}$ (and the vertices incident with those edges). Cycles are always countable subgraphs, because every edge on a circle contains a point that corresponds to a rational point on the unit circle. Moreover, every cycle is clearly 2-regular, and therefore either a finite cycle or a disjoint union of double rays.

In [5] we define the cycle space of a locally finite graph G essentially as the set of those sums $\sum_{i \in I} C_i$ of cycles C_i of G for which no edge of G occurs in C_i for infinitely many indices i (where $\sum_{i \in I} C_i$ denotes the subgraph of G consisting of those edges that lie in C_i for an odd number of indices i). In fact, in [5] we just considered the edge sets of cycles and their sums, rather than the cycles themselves. In the presence of infinite degrees however, we shall also have to take account of multiplicities of vertices if we want at least some spanning trees to exist whose fundamental cycles generate the cycle space. Indeed, let G be the graph obtained from two distinct vertices v and w by adding new vertices $x_1, x_2 \ldots$ and joining them to both v and w. Then the path $P = vx_1w$ is a well-defined sum of finite cycles according to the above definition (and hence an element of the cycle space), but there is no spanning tree T of G whose fundamental cycles sum to P: any such sum would consist of infinitely many fundamental cycles each containing v and w, and so the two edges of the path vTw would lie in infinitely many summands (contradiction). Hence there is no spanning tree of G for which the fundamental cycles generate its cycle space.

To overcome this problem we sharpen the requirements on the sums making up the cycle space, as follows. Call a family $(G_i)_{i \in I}$ of subgraphs of a graph G *thin* if no vertex of G lies in G_i for infinitely many i. Let the *sum* $\sum_{i \in I} G_i$ of this family be the subgraph of G consisting of all edges that lie in G_i for an odd number of indices i (and the vertices incident with these edges), and let the *cycle space* $\mathcal{C}(G)$ of G be the set of all sums of (thin families of) cycles. Then $\mathcal{C}(G)$ is closed under finite sums, and we shall see in Section 5 that it is even closed under infinite sums. Moreover, if G is finite then this definition is compatible with the standard one (except that we now consider subgraphs of G rather than edge sets). Similarly, if G is locally finite then our definition reduces to that given in [5].

We shall frequently use the following standard lemma about infinite graphs; the proof is not difficult and is included in [4, Lemma 1.2].

Lemma 2.4 *Let U be an infinite set of vertices in a connected graph G. Then G contains either a ray R with infinitely many disjoint U–R paths or a subdivided star with infinitely many leaves in U.* □

Let H be a subgraph of G. Then every end ω of H is a subset of a unique end ω' of G. The map $\pi_{HG} : \overline{H} \to \overline{G}$ which is the identity on H and sends every end ω of H to the end ω' of G containing it, is called the *canonical projection of $\overline{H}$ to $\overline{G}$*. Note that π_{HG} is continuous.

H is called *end-faithful* in G if π_{HG} maps the ends of H bijectively to the ends of G, i.e. if every end of G contains rays from exactly one end of H. H is *end-respecting* in G if π_{HG} is injective. Lemma 2.4 implies that end-respecting spanning subgraphs of locally finite graphs are end-faithful, but this is not true in general. We remark that normal spanning trees are end-faithful, with π^{-1} (as well as π) continuous [3].

Let us call H *separation-faithful* in G if a finite set $S \subseteq V(H)$ of vertices never separates two vertices of $H - S$ in H unless it also separates these two vertices in G. (Note that the converse always holds trivially.) In other words, for every finite $S \subseteq V(H)$ the components of $H - S$ are precisely the intersections of H with the components of $G - S$. If H is separation-faithful in G then, clearly, it is end-respecting. In fact, it is as close to end-faithful as its size allows, representing all the ends of G to which its vertices converge:

Lemma 2.5 *Let H be a separation-faithful subgraph of G. Then*

(i) *π_{HG} is a topological embedding, ie. π_{HG} is injective and π_{HG}^{-1} (as well as π_{HG}) is continuous;*

(ii) *$\pi_{HG}(\overline{H})$ is closed in $\overline{G}$.*

Proof. (i) is straightforward.

(ii) Let ω be a point in the closure of $\pi_{HG}(\overline{H})$ in $\overline{G}$; we wish to show that $\omega \in \pi_{HG}(\overline{H})$. Since π_{HG} is the identity on H, we may assume that ω is an end; let $R \subseteq G$ be a ray from ω. We shall construct a ray $Q \in \omega$ in H; then π_{HG} will map the end of Q in H to ω, as desired.

We start by constructing a countably infinite set $\mathcal{P}$ of disjoint H–R paths in G (possibly trivial); recall that an H–R path meets H and R only in its first and last vertex, respectively. This can be done inductively: having picked finitely many such paths, let S be the union of their vertex sets and recall that, since ω lies in the closure of $\pi_{HG}(\overline{H})$, the basic open neighbourhoods $\widehat{C}_G(S,\omega)$ of ω have a point (and hence a vertex) in H. We can then find our next H–R path in C.

Having completed the construction of $\mathcal{P}$, we let G' denote the union of R and all the paths in $\mathcal{P}$. Given vertices $x \in R$ and $y \in G'$, we say that y lies *above* x (and x *below* y) if y lies in the unique infinite component of $G' - x$. Similarly, if $x \in P \in \mathcal{P}$ and $y \in G'$, then y lies *above* x (and x *below* y) if y lies in the unique infinite component of $G' - P$. So only vertices lying on a common path in $\mathcal{P}$ are incomparable with respect to this relation; in particular, the vertices in $H \cap G'$ form an infinite increasing chain.

Pick any vertex $x_0 \in H \cap G'$, and set $Q_0 := \{x_0\}$. We shall now define paths $Q_1, Q_2, \ldots$ in H such that, for all $i \geq 1$, Q_i meets G' in its endvertices but in no other vertex, Q_i starts at the last vertex x_{i-1} of Q_{i-1}, is otherwise disjoint from $Q_0 \cup \ldots \cup Q_{i-1}$, and ends at a vertex $x_i \in G'$ above x_{i-1}. Then all the Q_i together will form a ray $Q \subseteq H$ which meets G' infinitely often, and which is therefore equivalent to R.

So let $i \geq 0$ and suppose that we have already constructed $Q_0, \ldots, Q_i$. Let S be the union of $V(Q_1 \cup \ldots \cup Q_i) \setminus \{x_i\}$ with the set of all vertices in $H \cap G'$ below x_i. By the properties assumed for $Q_1, \ldots, Q_i$ all of $S \cap G'$ lies below x_i, so x_i lies in the same component of $G' - S \subseteq G - S$ as the (infinitely many) vertices of $H \cap G'$ above it. Since H is separation-respecting, the same is true in $H - S$. So $H - S$ contains a path from x_i to another vertex of G' which we may choose as Q_{i+1}. □

When $H \subseteq G$ is separation-faithful, then Lemma 2.5 (i) says that we may think of $\overline{H}$ as a subspace of $\overline{G}$; in particular, circles in $\overline{H}$ remain circles in $\overline{G}$. (When H is not separation-faithful this will normally fail, as π_{HG} may identify distinct ends on an $\overline{H}$-circle into a single end of G.) Lemma 2.5 (ii), on the other hand, implies the converse: any circle in $\overline{G}$ whose edges all lie in H will already be a circle in $\overline{H}$, ie. H contains all the required ends too. Let us note these observations formally for later use:

Corollary 2.6 *Let H be a separation-faithful subgraph of G.*

(i) *If C is a circle in $\overline{H}$ then $\pi_{HG}(C)$ is a circle in $\overline{G}$.*

(ii) *If C is a circle in $\overline{G}$ and $C \cap G \subseteq H$, then $\pi_{HG}^{-1}(C)$ is a circle in $\overline{H}$.*

(iii) *The cycles of H are precisely the cycles of G that are subgraphs of H. In particular, $\mathcal{C}(H) \subseteq \mathcal{C}(G)$.*

Proof. (i) is immediate from Lemma 2.5 (i).

(ii) By Lemma 2.3, C is the closure of $C \cap G$ in $\overline{G}$. Since $C \cap G \subseteq H$, this implies that C lies in the closure of $H = \pi_{HG}(H)$ in $\overline{G}$, which by Lemma 2.5 (ii) is (contained in) $\pi_{HG}(\overline{H})$. So $C \subseteq \pi_{HG}(\overline{H})$, and thus $\pi_{HG}^{-1}(C)$ is well-defined; it is a circle in $\overline{H}$ by Lemma 2.5 (i).

(iii) The first assertion follows from (i) and (ii) together with the fact that π_{HG} is the identity on H and maps ends to ends. The second assertion follows. □

Lemma 2.7 *Every countable subgraph G' of G can be extended to a countable separation-faithful subgraph of G.*

Proof. Let us define a sequence $H_0 \subseteq H_1 \subseteq \ldots$ of countable subgraphs of G, as follows. Put $H_0 := G'$. Let H_{i+1} be a graph obtained from H_i by adding, for every finite set $S \subseteq V(H_i)$ and for every pair of distinct components D_1, D_2 of $H_i - S$ that are contained in a common component D of $G - S$, a D_1–D_2 path in D. Clearly if H_i is countable then so is H_{i+1}, and hence $H := \bigcup_{i \in \mathbb{N}} H_i$ too is countable.

Let us show that H is separation-faithful. Suppose on the contrary that for some finite $S \subseteq V(H)$ there are vertices $x_1, x_2 \in H - S$ that are separated by S in H but not in G. Let j be large enough that H_j contains both x_1 and x_2 as well as S. Then x_1, x_2 belong to distinct components D_1, D_2 of $H_j - S$ but to a common component D of $G - S$, so $D_1 \cup D_2 \subseteq D$. Hence by construction, $H_{j+1} \subseteq H$ contains an x_1–x_2 path avoiding S, contradicting the choice of x_1 and x_2. □

Since an infinite cycle C in a graph G is just a disjoint union of rays, it is never a cycle in itself, ie. in the graph C. A standard application of Corollary 2.6 and Lemma 2.7, however, will be that C is a cycle in some countable subgraph of G:

Lemma 2.8 *For every cycle C in a graph G there exists a countable subgraph H of G such that C is a cycle in H.*

Proof. Recall that cycles are countable subgraphs. By Lemma 2.7, G has a countable separation-faithful subgraph H such that C is a subgraph of H. By Corollary 2.6 (iii), C is also a cycle in H. □

When we consider spanning trees, the following observation from [5] shows that we shall want those to be end-respecting: any other spanning tree T would contain an infinite cycle, which – apart from being counterintuitive – could not be a sum of fundamental cycles. (Clearly, in any such sum each fundamental cycle present could be taken to occur exactly once, but then the sum would contain its chord and hence not lie in T.)

Lemma 2.9 *Let T be a spanning tree of a graph G, and assume that T contains no infinite cycle of G. Then T is end-respecting.*

Proof. Suppose T contains two rays from a common end ω of G which are inequivalent in T. Then these rays can be chosen so as to meet precisely in their common first vertex. Their union C and ω together then form a circle in $\overline{G}$, and so $C \subseteq T$ is a cycle. □

3 The generating theorem for cycles

In this section we characterize the spanning trees of a graph G whose fundamental cycles generate every cycle of G.

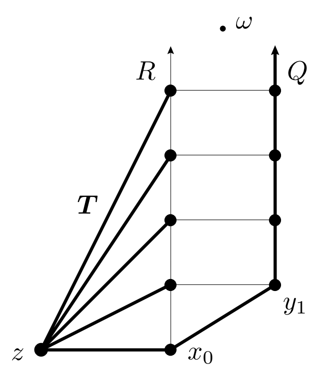

Figure 1: The infinite cycle $R \cup Q$ is not a sum of fundamental cycles

In [5] we showed that if G is locally finite, then these are precisely its end-respecting (equivalently: end-faithful) spanning trees. In general, however, this need not be true. Consider the graph G obtained from two rays $R = x_0x_1x_2\dots$ and $Q = x_0y_1y_2\dots$ that meet only in their first point x_0 by adding the edges x_iy_i for all $i \geq 1$, and adding a new vertex z joined to all the x_i (Fig. 1). Then R and Q belong to the same end ω of G. Thus $R \cup Q \cup \{\omega\}$ is a circle in $\overline{G}$, and so $R \cup Q$ is a cycle in G. But if T is the spanning tree of G consisting of Q together with all edges at z, then T is end-respecting (even end-faithful), but $R \cup Q$ is not a sum of fundamental cycles: since all these contain z, any sum of them would have to be finite.

The above example motivates the consideration of the following subclass of the end-respecting spanning trees.

Definition Given a graph G, let $\mathcal{T}(G)$ denote the class of all end-respecting spanning trees T of G which do not contain a subdivided infinite star S whose leaves lie on a ray $R \subseteq G$ such that G contains another ray R' which is equivalent to but disjoint from R.

Note that there are graphs G for which $\mathcal{T}(G)$ is empty; $K_{\aleph_1}$ and $K_{\aleph_0,\aleph_1}$ are obvious examples. On the other hand, using Lemma 2.2 one can easily show that $\mathcal{T}(G)$ contains every normal spanning tree of G. We do not know whether there are graphs G for which $\mathcal{T}(G)$ is non-empty but which have no normal spanning tree.

Theorem 3.1 *Let T be a spanning tree of G. Every cycle of G is the sum of fundamental cycles if and only if $T \in \mathcal{T}(G)$.*

For the proof of this theorem we first need some lemmas.

Lemma 3.2 *Given any spanning tree T of G, every finite cycle C of G is the sum of fundamental cycles. More precisely, C is equal to the sum Z of all the fundamental cycles C_e with $e \in E(C) \setminus E(T)$.*

Proof. Clearly $C + Z$ is a finite subgraph of T with all degrees even. Hence $C + Z = \emptyset$, i.e. $C = Z$. □

Lemma 3.3 *Let T be a spanning tree of G, let Z be a sum of fundamental cycles, and let $\mathcal{D}$ be a set of finite cycles in $Z \cup T$. If no two elements of $\mathcal{D}$ share an edge outside T, then $\mathcal{D}$ is thin.*

Proof. Suppose that x is a vertex that lies on infinitely many cycles $D \in \mathcal{D}$. By Lemma 3.2, each of these D is a sum of fundamental cycles C_e with $e \in E(D) \setminus E(T)$, so x lies on some C_e with $e \in E(D) \setminus E(T)$. By assumption these edges e differ for different D, so x lies on infinitely many C_e. As each e lies in Z, all these C_e are among the fundamental cycles whose sum is Z (indeed, Z must be the sum of the fundamental cycles C_e with $e \in E(Z) \setminus E(T)$), which contradicts the definition of sum. □

Lemma 3.4 *Let A be an arc in $\overline{G}$, and let $x \neq y$ be vertices on A. Let X be a closed subset of $\overline{G}$ which avoids the subarc of A between x and y. Then G contains an x–y path P with $P \cap X = \emptyset$.*

Proof. Let A' be the subarc of A between x and y. Choose a cover $\mathcal{N}$ of A' by basic open sets of $\overline{G}$ each avoiding X. As A' is compact, $\mathcal{N}$ contains a finite subcover of A', $\{N_1, \ldots, N_k\}$ say, where we may assume that $N_\ell \cap A' \neq \emptyset$ for all ℓ.

Let us show that $H := (N_1 \cup \cdots \cup N_k) \cap G$ is a connected subspace of G. If not, then H is the union of two disjoint non-empty open subsets H_1 and H_2 of H. Since each N_ℓ is a basic open set, $N_\ell \cap G$ is connected and hence lies in either H_1 or H_2. Let U_1 be the union of all N_ℓ with $N_\ell \cap G \subseteq H_1$, and define U_2 similarly. Since two N_ℓ cannot share an end if their intersections with G are disjoint, U_1 and U_2 are disjoint. Thus A' is the union of the two disjoint non-empty open sets $A' \cap U_1$ and $A' \cap U_2$, contradicting its connectedness.

So H is connected. Lemma 2.1 together with the fact that H contains both x and y imply that H also contains a (graph-theoretical) path P between these two vertices. Clearly, P is as required. $\square$

An *orientation* of an arc A is a linear ordering of its points which is induced by a homeomorphism $\sigma : [0,1] \to A$ (i.e. if $a, b \in A$ then $a < b$ if $\sigma^{-1}(a) < \sigma^{-1}(b)$ in $[0,1]$). Given an oriented arc $\vec{A}$ and $a \in A$, we will refer to the points $b \in A$ with $b < a$ as the points *left* of a, and analogously we will speak of points to the *right* of a. We write $a\vec{A}$ for the (oriented) subarc of A consisting of all the points $a' \geq a$, and define $\vec{A}a$ and $a\vec{A}b$ analogously. A sequence $(e_i)_{i=1}^{\infty}$ of distinct edges or vertices on A is *monotone* if there is an orientation of A such that each e_i lies *between* e_{i-1} and e_{i+1}, i.e. on the right of e_{i-1} and on the left of e_{i+1}. A sequence $(e_i)_{i=1}^{\infty}$ of distinct edges or vertices on a circle C is *monotone* if there is a subarc A of C containing each e_i and $(e_i)_{i=1}^{\infty}$ is monotone on A. An *orientation* of C is a choice of one of the two orientations of every arc $A \subseteq C$ such that all these orientations are compatible on their intersections. Given an oriented circle $\vec{C}$ and $a, b \in C$ with $a \neq b$ we define $a\vec{C}b$ to be the (oriented) subarc of C between a and b.

Lemma 3.5 *Let A be an arc in $\overline{G}$. Let $(e_i)_{i=1}^{\infty}$ and $(f_i)_{i=1}^{\infty}$ be monotone sequences of distinct edges on A converging from different sides to an end ω of G lying on A. Then ω contains two disjoint rays R and R' such that R contains every e_i while R' contains every f_i.*

Proof. First fix an orientation of A. We may assume that $(e_i)_{i=1}^{\infty}$ converges to ω from the left, and $(f_i)_{i=1}^{\infty}$ converges to ω from the right. Let $e_i =: x_i^1 x_i^2$

and $f_i =: y_i^1 y_i^2$ where x_i^1 lies on the left of x_i^2 and y_i^1 lies on the right of y_i^2. Let $A_i := x_i^2 \vec{A} x_{i+1}^1$ and $A_i' := \vec{A} x_1^1 \cup x_{i+1}^2 \vec{A}$, and let $B_i := y_{i+1}^1 \vec{A} y_i^2$ and $B_i' := y_1^1 \vec{A} \cup \vec{A} y_{i+1}^2$.

We will construct the rays R and R' inductively, extending in each step the initial segments of R and R' already defined. Thus suppose that for some $i \geq 0$ we have constructed finite disjoint paths R_i and R_i' which are empty if $i = 0$, and for $i > 0$ are such that R_i joins x_1^2 to x_{i+1}^1, contains each e_j with $1 < j \leq i$ and avoids A_i', while R_i' joins y_1^2 to y_{i+1}^1, contains each f_j with $1 < j \leq i$ and avoids B_i'.

Let us now extend R_i and R_i'. By Lemma 3.4 there is an x_{i+1}^2–x_{i+2}^1 path P in G which avoids the closed set $R_i \cup R_i' \cup A_{i+1}'$. Put $R_{i+1} := R_i e_{i+1} P$. Applying Lemma 3.4 again, we find a y_2^{i+1}–y_1^{i+2} path P' which avoids $R_{i+1} \cup R_i' \cup B_{i+1}'$. Put $R_{i+1}' := R_i' f_{i+1} P'$. Continuing inductively, we obtain rays $R := \bigcup_{i=1}^{\infty} R_i$ and $R' := \bigcup_{i=1}^{\infty} R_i'$. But then $e_1 R$ and $f_1 R'$ are as required. □

Lemma 3.6 *Let T be a spanning tree of G, and let T_1, T_2 be subtrees of T with finite intersection. Suppose that G has an end ω which, for each $i = 1, 2$, contains disjoint rays R_i and R_i' such that R_i has infinitely many vertices in T_i. Then $T \notin \mathcal{T}(G)$.*

Proof. For $i = 1, 2$, apply Lemma 2.4 to T_i with $U := V(R_i \cap T_i)$. If the lemma returns a star in one of the T_i then $T \notin \mathcal{T}(G)$ by definition of $\mathcal{T}(G)$. But if it returns a ray in each T_i then both these rays lie in ω, and so T is not end-respecting. Thus again $T \notin \mathcal{T}(G)$. □

We will also need the following lemma from elementary topology [10, p. 208]. A continuous image of $[0, 1]$ in a topological space X is a (topological) *path* in X; the images of 0 and 1 are its *endpoints*.

Lemma 3.7 *Every path with distinct endpoints x, y in a Hausdorff space X contains an arc in X between x and y.* □

Proof of Theorem 3.1. To prove the forward implication, suppose that $T \notin \mathcal{T}(G)$. By Lemma 2.9 and the remark preceding it we may assume that T is end-respecting. Thus there are disjoint equivalent rays R and R' in G such that T contains a subdivision S of an infinite star whose leaves lie on R. Clearly, we may assume that R meets S only in its leaves. Let ω be the end of G containing R and R'. Let $P = x \dots x'$ be an R–R' path in G. Let C' be the circle in $\overline{G}$ consisting of ω together with P, xR and $x'R'$. Let C

be the cycle of C'. Thus $C = P \cup xR \cup x'R'$. Let $\mathcal{D}$ be the (infinite) set of all finite cycles which consist of a finite subpath of xR between two consecutive leaves of S on xR together with the path in S joining these leaves. Then $\mathcal{D}$ is not thin, since the centre of S lies in all cycles in $\mathcal{D}$. Lemma 3.3 now implies that C cannot be a sum of fundamental cycles, as required.

To prove the converse implication, we now assume that $T \in \mathcal{T}(G)$. Let C be a cycle of G; we shall prove that C is the sum of all the fundamental cycles C_e of T with $e \in C$. Let $\mathcal{C}$ denote the set of these C_e. Let C' be the defining circle of C, let C'' be the unit circle, and let $\sigma : C'' \to C'$ be a homeomorphism.

We first show that $\mathcal{C}$ is a thin family. Suppose not, and let x be a vertex that lies on C_e for infinitely many chords e of T on C. Since C' is compact, these edges e have an accumulation point ω on C' (which must be an end), and we may choose a monotone sequence $e_1, e_2, \dots$ from among these edges that converges to ω. Since $x \in C_{e_i}$, the endvertices of e_i never lie in the same component of $T - x$. Partitioning the components of $T - x$ suitably into two sets, we may write T as the union of two subtrees T_1 and T_2 that meet precisely in x and are joined by infinitely many e_i. Applying Lemma 3.5 to a suitable subarc of C' containing all the e_i as well as a monotone sequence of edges on C' converging to ω from the other side, we obtain disjoint rays R and R' both belonging to ω and such that R contains every e_i. Then R meets both T_1 and T_2 infinitely often, and we may apply Lemma 3.6 with $R_1 := R =: R_2$ and $R'_1 := R' =: R'_2$ to conclude that $T \notin \mathcal{T}(G)$, contrary to our assumption.

It remains to prove that the cycles in $\mathcal{C}$ sum to C. We thus have to show that an edge f of G lies on an odd number of the cycles in $\mathcal{C}$ if and only if $f \in C$. This is clear when f is a chord of T (and C_f is a fundamental cycle), so we assume that $f \in T$. Let G_1 and G_2 be the subgraphs of G induced by the components of $T - f$, and let E_f be the set of all G_1–G_2 edges of G (including f). Note that the edges $e \neq f$ in E_f are precisely the chords e of T with $f \in C_e$. As $\mathcal{C}$ is thin, C contains only finitely many edges from E_f.

Let us show that the number of edges of C in E_f is even. Since σ is a homeomorphism, $C'' \setminus \sigma^{-1}(\mathring{E}_f \cap C)$ consists of finitely many closed intervals, $I_1, \dots, I_k$ say. Since each $\sigma(I_i) \subseteq C'$ is path-connected, it suffices to show that G_1 and G_2 belong to different path-components X_1 and X_2 of $\overline{G} \setminus \mathring{E}_f$: then each $\sigma(I_i)$ lies inside either X_1 or X_2, and thus $E(C) \cap E_f$ is even. Suppose then that G_1 and G_2 are contained in the same path-component of $\overline{G} \setminus \mathring{E}_f$. By Lemma 3.7, there is an arc A in $\overline{G} \setminus \mathring{E}_f$ from a vertex of G_1 to one in G_2. Let ω be the supremum of the points on A that lie in G_1; this can only be an end. Choose monotone edge sequences $(e_i)_{i=1}^{\infty}$ and $(f_i)_{i=1}^{\infty}$ on A with all e_i in G_1 and all f_i in G_2, and so that $(e_i)_{i=1}^{\infty}$ and $(f_i)_{i=1}^{\infty}$ converge to ω from different sides. Apply Lemma 3.5 to obtain disjoint rays R and R' in ω such that R contains every e_i while R' contains every f_i. Now Lemma 3.6 applied with $R_1 := R =: R'_2$ and $R_2 := R' =: R'_1$ implies that $T \notin \mathcal{T}(G)$, a contradiction.

So we have proved that C contains an even number of edges from E_f. As $f \in E_f$, this means that $f \in C$ if and only if C contains an odd number of the edges $e \neq f$ from E_f, which it does if and only if f lies on an odd number of fundamental cycles $C_e \in \mathcal{C}$. □

4 Generating arbitrary elements of the cycle space

In this section we characterize the spanning trees whose fundamental cycles generate not only each individual cycle but the entire cycle space of an arbitrary graph. It turns out that these include all normal spanning trees. We shall need this fact in the proof of our characterization theorem below, so let us prove it first:

Lemma 4.1 *Let G be a graph with a normal spanning tree T. Then every element Z of the cycle space of G is the sum of fundamental cycles.*

Proof. Write Z as the sum $\sum_{i \in I} Z_i$ of cycles of G. Since $\mathcal{T}(G)$ contains the normal spanning tree T, Theorem 3.1 implies that each Z_i is a sum $\sum_{j \in J_i} C_i^j$ of fundamental cycles. We may assume that the C_i^j are distinct for different $j \in J_i$. To prove the lemma, it suffices to show that the family $\mathcal{C} := (C_i^j)_{i \in I, j \in J_i}$ is thin, since then clearly Z is the sum of all the cycles in $\mathcal{C}$. So suppose that $\mathcal{C}$ is not thin. Then there is a vertex v which lies in the fundamental cycles C_i^j for an infinite set J of pairs (i, j). Since T is normal, every vertex set of a fundamental cycle C_e is a chain in T, its minimum and maximum being joined by e. Thus choosing v minimal in T and possibly discarding finitely many pairs from J, we may assume that v is the lowest vertex (in T) of each C_i^j with $(i, j) \in J$ and thus incident with its chord e_i^j. As C_i^j is the only cycle in the family $(C_i^j)_{j \in J_i}$ that contains e_i^j and this family sums to Z_i, we have $v \in e_i^j \in Z_i$ for all $(i, j) \in J$. But each Z_i has only finitely many summands C_i^j containing v, so $v \in Z_i$ for infinitely many i. Thus $(Z_i)_{i \in I}$ is not thin, contradicting the fact that $Z = \sum_{i \in I} Z_i$. □

We remark that Lemma 4.1 does not extend to arbitrary spanning trees in $\mathcal{T}(G)$. For example, consider the graph G obtained from infinitely many disjoint finite cycles $C_1, C_2, \ldots$ by adding a new vertex s and joining it to two vertices of each C_i. Let T be a spanning tree of G containing all the edges of G incident with s. Then $T \in \mathcal{T}(G)$. But as each fundamental cycle contains s, the element $Z = \sum_{i=1}^{\infty} C_i$ of the cycle space of G is not a sum of fundamental cycles.

Let us then determine the subclass $\mathcal{T}'(G) \subseteq \mathcal{T}(G)$ of those spanning trees of G whose fundamental cycles generate all of $\mathcal{C}(G)$. A *comb* C with *back* R is obtained from a ray R and a sequence $x_1, x_2, \ldots$ of distinct vertices (to be called the *teeth* of C) by adding for each $i = 1, 2, \ldots$ a (possibly trivial) x_i–R path P_i so that all the P_i are disjoint.

Definition Let $\mathcal{T}'(G)$ be the class of all spanning trees $T \in \mathcal{T}(G)$ such that G does not contain infinitely many disjoint finite cycles $C_1, C_2, \ldots$ for which one of the following conditions holds (Fig. 2):

- T contains two subdivided infinite stars S_1 and S_2 such that S_1 and S_2 meet at most in the centre of S_1 which is then also the centre of S_2, and such that each C_i contains a leaf of both S_1 and S_2 $(i = 1, 2, \ldots)$.
- T contains a subdivided infinite star S and a comb C such that S and C are disjoint and each C_i contains both a leaf of S and a tooth of C $(i = 1, 2, \ldots)$.

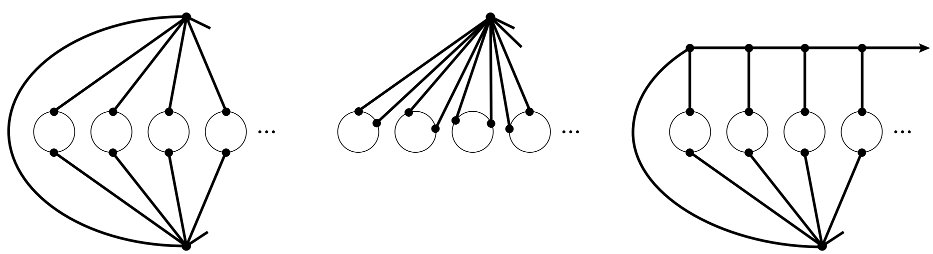

Figure 2: The additional forbidden configurations for $\mathcal{T}'(G)$

As before, one can easily show using Lemma 2.2 that $\mathcal{T}'(G)$ contains every normal spanning tree of G.

Theorem 4.2 *Let T be a spanning tree of a graph G. Every element of the cycle space of G is a sum of fundamental cycles if and only if $T \in \mathcal{T}'(G)$.*

For the proof of this theorem we again need a few lemmas. First recall the following basic fact from point-set topology (see e.g. [1, Thm. 3.7]):

Lemma 4.3 *Every continuous injective map from a compact space X to a Hausdorff space Y is a topological embedding, i.e. a homeomorphism between X and its image in Y under the subspace topology.* □

Lemma 4.4 *Let $H_1 \subseteq H_2$ be subgraphs of G, and let C be a cycle in H_1. If C is a cycle in G, then it is also a cycle in H_2.*

Proof. Let C' be the defining circle of C in $\overline{H}_1$. We show that the restriction to C' of the canonical embedding $\pi_{H_1 H_2}$ is injective; then by Lemma 4.3 it is a topological embedding (since $\pi_{H_1 H_2}$ is continuous), and so $C = \pi_{H_1 H_2}(C') \cap H_2$ will be a cycle in H_2.

Note first that $\pi_{H_1 G}$ maps C' onto the defining circle C'' of C in $\overline{G}$: since $\pi_{H_1 G}(C')$ is compact (and hence closed) and contains C as a dense subset, it is the closure of C in $\overline{G}$, which we know to be C''.

Now if $\pi_{H_1 H_2}$ is not injective on C' then neither is $\pi_{H_1 G} = \pi_{H_2 G} \circ \pi_{H_1 H_2}$, so there are two ends $\omega_1, \omega_2 \in C'$ with $\pi_{H_1 G}(\omega_1) = \pi_{H_1 G}(\omega_2)$. Pick $x, y \in C$ so that ω_1, ω_2 lie in distinct path-components of $C' \setminus \{x, y\}$. Then $\pi_{H_1 G}(C' \setminus \{x, y\}) = C'' \setminus \{x, y\}$ is path-connected, contradicting the fact that removing any two distinct points from a circle makes it path-disconnected. □

Lemma 4.5 *Let T be a spanning tree of G, and let $C_1, C_2, \ldots \subseteq G$ be disjoint finite cycles. From each C_i pick an edge e_i not on T. If G has a vertex x that lies on each of the fundamental cycles C_{e_i}, then $T \notin \mathcal{T}'(G)$.*

Proof. As $x \in C_{e_i}$, each e_i has its endvertices in two different components of $T - x$. Partitioning these components suitably into two sets, we may write T as the union of two subtrees T_1 and T_2 that meet precisely in x and are joined by infinitely many e_i. Applying Lemma 2.4 to T_1 with U the set of endvertices of these e_i in T_1, we obtain an infinite set $I \subseteq \mathbb{N}$ and either a ray in T_1 joined to all the e_i with $i \in I$ by disjoint paths in T_1, or else a subdivided star in T_1 whose leaves are precisely the endvertices of the e_i with $i \in I$ in T_1. Now apply Lemma 2.4 to T_2 with U the set of endvertices of these e_i $(i \in I)$ in T_2 to obtain an infinite set $I' \subseteq I$ and either a ray or a subdivided star in T_2. If both applications of the lemma return a ray then these rays are equivalent, and so T is not end-respecting. If both return a star, then these stars can be chosen so as to meet at most in their common centre (which then must be x). As $e_i \in C_i$, each C_i with $i \in I'$ contains leaves of both stars. So these stars satisfy the first condition in the definition of $\mathcal{T}'(G)$. Similarly, if the lemma returns a ray and a star, then they satisfy the second condition in the definition of $\mathcal{T}'(G)$. Thus in each case we have shown that $T \notin \mathcal{T}'(G)$, as desired. □

Proof of Theorem 4.2. To prove the forward implication, suppose that $T \notin \mathcal{T}'(G)$. By Theorem 3.1 we may assume that $T \in \mathcal{T}(G)$. Thus there are disjoint finite cycles $C_1, C_2, \ldots$ in G satisfying one of the two conditions in the definition of $\mathcal{T}'(G)$. We consider only the case that T contains two subdivided infinite stars S_1 and S_2 (which are either disjoint or meet only in their common centre) such that each C_i meets both S_1 and S_2; the other case is similar. We may assume that $C_1 \cup C_2 \cup \ldots$ avoids the path $P \subseteq T$ joining the centre of S_1 to that of S_2. On each C_i choose an S_1–S_2 path $P_i = x_i \ldots y_i$. Since C_i is disjoint from P, the x_i–y_i path in T forms a finite cycle together with P_i. Let $\mathcal{D}$ denote the set of all these cycles, one for each i. Then $\mathcal{D}$ is not thin, as every cycle in $\mathcal{D}$ contains the centre of S_1. Thus Lemma 3.3 implies that the element $Z = \sum_{i=1}^{\infty} C_i$ of the cycle space of G cannot be the sum of fundamental cycles, as desired.

To prove the converse implication, suppose that $T \in \mathcal{T}'(G)$, and let Z be an element of the cycle space of G. Write Z as the sum $\sum_{i \in I} Z_i$ of cycles Z_i. By Theorem 3.1, each Z_i is the sum of a thin family $\mathcal{C}_i = (C_i^j)_{j \in J_i}$ of (distinct) fundamental cycles. It suffices to show that $\mathcal{C} := (C_i^j)_{i \in I, j \in J_i}$ is a thin family: then clearly Z is the sum of all the cycles in $\mathcal{C}$.

Suppose that $\mathcal{C}$ is not thin. Then some vertex x lies on infinitely many cycles in $\mathcal{C}$. Since every family $\mathcal{C}_i$ is thin, there exists an infinite set $I' \subseteq I$ such that for every $i \in I'$ the vertex x lies on some cycle in $\mathcal{C}_i$. Denoting the defining chord of this (fundamental) cycle by e_i, we thus have $x \in C_{e_i} \in \mathcal{C}_i$ for every $i \in I'$.

As the fundamental cycles in $\mathcal{C}_i$ are distinct, their defining chords do not cancel in the sum $\sum_{C \in \mathcal{C}_i} C = Z_i$, so $e_i \in Z_i$ for every i. On the other hand as the family $(Z_i)_{i \in I}$ is thin, we have $e_i \in Z_k$ for only finitely many k. In particular, $e_i \neq e_k$ for all but finitely many k. Conversely, Z_k contains only finitely many e_i (since $\mathcal{C}_k$ is thin and every C_{e_i} contains x), so $Z_k \ni e_i$ for only finitely many i. Replacing I' with an appropriate infinite subset if necessary, we may therefore assume that $e_i \in Z_k$ if and only if $i = k$ (for all $i, k \in I'$), and further that I' is countable.

For $Z' := \sum_{i \in I'} Z_i$ the above implies that $e_i \in Z'$ for all $i \in I'$. Moreover, Lemmas 2.8 and 4.4 imply that Z' lies in the cycle space of a countable subgraph H of G. Since every countable connected graph has a normal spanning tree, Lemma 4.1 thus implies that Z' is a sum of a thin family

$\mathcal{C}'$ of finite cycles: of fundamental cycles of normal spanning trees of the components of H. As every e_i lies in Z' and hence in some cycle of $\mathcal{C}'$, and since each of these cycles meets only finitely many others, $\mathcal{C}'$ has an infinite subfamily of disjoint cycles each containing an edge e_i with $i \in I'$. Lemma 4.5 now implies that $T \notin \mathcal{T}'(G)$, contradicting our assumption.

□

5 Cycle decompositions and cycle-cut orthogonality

In this section we establish infinite analogues of two further properties of finite cycle spaces, properties that make no reference to spanning trees: the fact that every element of the cycle space of a finite graph is an edge-disjoint union of cycles (Theorem 5.2), and that the cycle space consists of precisely those (edge sets of) subgraphs that are 'orthogonal' to every cut (Theorem 5.4).

The basic idea for the proof of Theorem 5.2 is the same as in the finite case: given $Z \in \mathcal{C}(G)$, we shall find a single cycle $C \subseteq Z$ in G and then iterate with $Z - E(C)$, continuing until the cycles deleted from Z have exhausted it. As in the finite case, none of the cycles from the constituent sum of Z need be a subgraph of Z, so finding C is non-trivial. But while for finite Z we can find C greedily inside Z (using the fact that all degrees in Z are at least 2), this need not be possible when Z is infinite: a maximal path in Z may be any double ray not defining a cycle, and it is then not clear how to extend this double ray beyond its ends to a circle giving rise to the desired cycle C.

Our main lemma for the proof of Theorem 5.2 thus deals with finding C, and it does so in a countable subgraph H of G. Finding the right H in which to do this will cause a few (managable) complications later on, but the key advantage is that H, being countable, will have a normal spanning tree T. We may then write any $Z \in \mathcal{C}(H)$ as a sum of *finite* cycles (namely, of fundamental cycles with respect to T; cf. Lemma 4.1), which will make standard compactness arguments available for the construction of C.

Lemma 5.1 *Let H be a countable subgraph of G, let $Z \in \mathcal{C}(H)$, and let $e = vw \in E(Z)$. Then $\overline{H}$ contains a topological path P from v to w such that $P \cap H \subseteq Z - e$.*

Proof. As H is countable, it has a normal spanning tree. Thus Lemma 4.1 implies that Z can be written as $Z = \sum_{i=1}^{\infty} C_i$, where the C_i are finite cycles in H forming a thin family. Let $H' := \bigcup_{i=1}^{\infty} C_i$. Replacing Z with the sum Z' of those C_i that lie in the component of H' containing e, we may assume that H' is connected. (Indeed, $Z' \in \mathcal{C}(H)$ and $e \in Z' \subseteq Z$; hence a proof of the lemma for Z' implies it for Z.) Since the family $(C_i)_{i=1}^{\infty}$ is thin, H' is locally finite. Put $Z_i := \sum_{j=1}^{i} C_j$. As $e \in Z$, there exists an $i_0 > 0$ such that $e \in Z_i$ for all $i \geq i_0$. Furthermore, each Z_i is finite and hence an edge-disjoint union of finite cycles in H'. Fix such a set of finite cycles for every $i \geq i_0$, and let D_i denote the cycle containing e. Let P_i be the finite path $D_i - e$, and orient it from v to w.

Let $e_1, e_2, \ldots$ be an enumeration of the edges in $E(H') \setminus \{e\}$. Let us define a sequence $X_0 \subseteq X_1 \subseteq \ldots$ of finite subsets of $E(H') \setminus \{e\}$ and a sequence $I_0 \supseteq I_1 \supseteq \ldots$ of infinite subsets of $\mathbb{N}$ so that the following holds for all $i = 0, 1 \ldots$:

> $X_i = \{e_1, \ldots, e_i\} \cap E(P_j)$ *for all* $j \in I_i$, *and all these* P_j *induce the same linear ordering on* X_i *and the same orientation on the edges in* X_i. $(*)$

To do this, we begin with $X_0 = \emptyset$ and $I_0 = \{i \in \mathbb{N} \mid i \geq i_0\}$. For every $i \geq 0$ in turn, we then check whether $e_{i+1} \in P_j$ for infinitely many $j \in I_i$. If so, we put $X_{i+1} := X_i \cup \{e_{i+1}\}$ and choose $I_{i+1} \subseteq I_i$ so as to satisfy $(*)$ for $i+1$; if not, we let $X_{i+1} := X_i$ and put $I_{i+1} := \{j \in I_i \mid e_{i+1} \notin P_j\}$ (in which case $I_i \setminus I_{i+1}$ is finite, and $(*)$ again holds for $i+1$). Finally, let $X := \bigcup_{i=0}^{\infty} X_i$, and write $\dot{X}$ for the subgraph of H consisting of the edges in X and their incident vertices.

The set X is linearly ordered as follows. Given $f, f' \in X$, consider the least index i such that $f, f' \in X_i$. If f precedes f' (say) on one P_j with $j \in I_i$ then it does so on every such P_j, and hence in particular on every P_j with $j \in I_k$ and $k > i$ (since $I_k \subseteq I_i$). Similarly, every edge $f \in X$ has a unique orientation, its common orientation on every P_j with $j \in I_i$ and i large enough that $f \in X_i$.

Let us show that $\dot{X} \subseteq Z - e$. Given an edge $f \in X$, we have $f \in P_j \subseteq Z_j - e$ for infinitely many j; indeed, by $(*)$ this holds for all $j \in I_i$ with i large enough that $f \in X_i$. But then $f \in Z_j$ for all large enough j (because f lies on only finitely many C_i), and hence also $f \in Z$.

Using the local finiteness of H', it is in fact easy to show that $\dot{X} + e$ is a 2-regular subgraph of Z, in which two edges of X are adjacent if and only if they are adjacent elements in the linear ordering on X. Indeed, given a vertex $u \in \dot{X}$ choose k large enough that every edge of H' incident with u precedes e_k in the enumeration of all the edges e_i, and pick $j \in I_k$. Then the edges at u in $\dot{X}$ are precisely the edges at u in X_k, which by $(*)$ are precisely the edges at u in P_j. If $u \in \{v, w\}$ there is one such edge; otherwise there are two.

If X is finite, then $\dot{X}$ is a v–w path in $Z - e$, and thus $\dot{X}$ is a topological path P as sought in the lemma. So let us assume that X is infinite. Then $\dot{X}$ is a disjoint union of two rays R_v and R_w starting at v and w, respectively, and possibly some further double rays. We will show that the closure of $\dot{X}$ in $\overline{H}$ is a topological path P as desired.

Assign to R_v a half-open subinterval J_{R_v} of $[0, 1]$ containing 0, to R_w a half-open subinterval J_{R_w} containing 1, and to each double ray $D \subseteq \dot{X}$ an open subinterval J_D, in such a way that all these intervals are disjoint, their order on $[0, 1]$ (oriented from 0 to 1) reflects the order of their corresponding rays and double rays induced by the linear ordering on X, and so that $[0, 1]$ is the closure of the union U of these subintervals. (Since $\dot{X}$ contains only countably many double rays, this can be done in at most ω steps.) Let $\sigma : [0, 1] \to \overline{H}$ map each interval J_Q continuously and bijectively onto its ray or double ray Q so that the order of the edges of Q in X reflects that induced by σ. Thus in particular $\sigma(0) = v$ and $\sigma(1) = w$. In what follows we will show that we can extend σ to a continuous map from $[0, 1]$ to $\overline{H}$ by mapping the points in $[0, 1] \setminus U$ to suitable ends of H. The image of $[0, 1]$ will then be a path P as desired.

So let x be a point in $[0, 1] \setminus U$. Choose a sequence $(u_i)_{i=1}^{\infty}$ of vertices of $\dot{X}$ so that the sequence $(\sigma^{-1}(u_i))_{i=1}^{\infty}$ is monotone in $[0, 1]$ and converges to x. Since H' is connected and locally finite, we may apply Lemma 2.4 to find a ray $Q_x \subseteq H'$ such that H' contains infinitely many disjoint Q_x–$\{u_i \mid i \in \mathbb{N}\}$ paths. Let ω_x be the end of H containing Q_x, and extend σ by setting $\sigma(x) := \omega_x$. (We remark that although formally ω_x depends on the choice of $(u_i)_{i=1}^{\infty}$, this is not in fact the case. However, we shall not need this below.)

We have to prove that $\sigma : [0, 1] \to \overline{H}$ is continuous. Clearly, σ is continuous in points of U. So let $x \in [0, 1] \setminus U$, and let N be a basic open neighbourhood of ω_x in $\overline{H}$. Then N is of the form $\widehat{D}$ for some component D of $H - S$ with $S \subseteq V(H)$ finite. We have to show that there is an open neighbourhood O of x in $[0, 1]$ such that $\sigma(O) \subseteq \widehat{D}$.

We will first show that there are points $a \neq b$ in $[0, 1]$ such that $x \in (a, b)$ and either $\sigma(a, x) \cap \dot{X} \subseteq D$ or $\sigma(a, x) \cap \dot{X} \cap D = \emptyset$, and such that the analogous assertion holds for (x, b). Let $k := |S|$, and suppose there is no such point a (say). Then we can find a monotone sequence $f_1, \ldots, f_{k+2}$ of $k + 2$ distinct edges in X lying alternately inside and outside of D (and having no incident vertex in S). As the sequence $f_1, \ldots, f_{k+2}$ is monotone in the ordering on X (and this ordering is well defined), every path P_j with $j \in I_i$ and i large enough that $f_1, \ldots, f_{k+2} \in X_i$ contains all these edges in this order. But then P_j meets S in at least $k+1$ vertices, a contradiction. Hence there are points a and b as required.

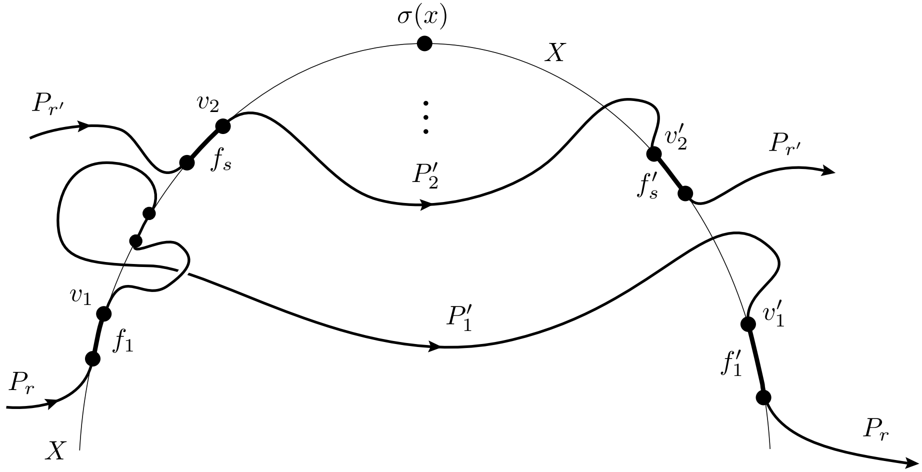

Figure 3: Constructing the paths P'_i

Let us now show that either $\sigma(a,b) \cap \dot{X} \subseteq D$ or $\sigma(a,b) \cap \dot{X} \cap D = \emptyset$. This will follow from the choice of a and b if there are sequences $(v_i)_{i=1}^{\infty}$ and $(v'_i)_{i=1}^{\infty}$ of distinct vertices of $\dot{X}$ such that $(\sigma^{-1}(v_i))_{i=1}^{\infty}$ is monotone and converges to x from the left while $(\sigma^{-1}(v'_i))_{i=1}^{\infty}$ is monotone and converges to x from the right, and such that H contains infinitely many disjoint paths $P'_i = v_i \dots v'_i$. We will construct such paths inductively (Fig. 3). Let $(f_i)_{i=1}^{\infty}$ and $(f'_i)_{i=1}^{\infty}$ be monotone sequences of distinct edges of X such that $(\sigma^{-1}(f_i))_{i=1}^{\infty}$ converges to x from the left while $(\sigma^{-1}(f'_i))_{i=1}^{\infty}$ converges to x from the right, and such that f_{i+1} succeeds both f_i and f'_{i+1} in the enumeration $e_1, e_2, \dots$ of $E(H') \setminus \{e\}$, and f'_{i+1} succeeds f'_i in this enumeration (for all $i \geq 1$). Let k be such that $f_1 = e_k$, and pick $r \in I_k$. Then $f_1, f'_1 \in P_r$: if $i < k$ is such that $e_i = f'_1$, then $r \in I_k \subseteq I_i$ and hence $f'_1 = e_i \in X_i \subseteq E(P_r)$ by $(*)$. Moreover, since f_1 lies to the left of f'_1 in X, it precedes f'_1 on P_r. Let v_1 be the last vertex of f_1 on P_r, and let v'_1 be the first vertex of f'_1 on P_r. Put $P'_1 := v_1 P_r v'_1$. Now let $s > 1$ be such that f_s succeeds every edge of P'_1 in the sequence $e_1, e_2, \dots$, and such that no edge of $E(P'_1) \cap X$ lies between f_s and f'_s in X. Let k' be such that $f_s = e_{k'}$, and pick $r' \in I_{k'}$. Then $f_s, f'_s \in P_{r'}$, and f_s precedes f'_s on $P_{r'}$. Let v_2 be the last vertex of f_s on $P_{r'}$, and let v'_2 be the first vertex of f'_s on $P_{r'}$. Put $P'_2 := v_2 P_{r'} v'_2$. Since $e_{k'}$ succeeds every edge from $E(P'_1) \setminus X$ in the enumeration of the e_i, condition $(*)$ implies that $P_{r'}$ (and hence P'_2) has no edge in $E(P'_1) \setminus X$. And P'_2 has no edge in $E(P'_1) \cap X$, because none of those edges lies between f_s and f'_s in X: since $e_{k'}$ equals or succeeds f_s, f'_s and every edge from $E(P'_1) \cap X$ in the enumeration of the e_i, the position of any such edge on $P_{r'}$ relative to f_s and f'_s would be the same as in X, i.e. it would precede f_s or succeed f'_s on $P_{r'}$ and hence not lie on P'_2. Thus P'_1 and P'_2 are edge-disjoint. Continuing inductively, we obtain infinitely many edge-disjoint paths $P'_i = v_i \dots v'_i$, one for every $i \in \mathbb{N}$. As all these paths lie in the locally finite graph H', infinitely many of them are disjoint, as desired. Thus we have shown that either $\sigma(a,b) \cap \dot{X} \subseteq D$ or $\sigma(a,b) \cap \dot{X} \cap D = \emptyset$.

By definition, ω_x contains the ray Q_x, and Q_x was defined in such a way that there is a sequence $(u_i)_{i=1}^{\infty}$ of distinct vertices in $\dot{X}$ such that H' contains infinitely many disjoint Q_x–$\{u_i \mid i \in \mathbb{N}\}$ paths, and where $(\sigma^{-1}(u_i))_{i=1}^{\infty}$ converges to x. Then all but finitely many of the points $\sigma^{-1}(u_i)$ lie in (a,b). Since $\sigma(x) = \omega_x \in \widehat{D}$, it follows that $\sigma(a,b) \cap \dot{X} \subseteq D$. Now let $y \in (a,b)$ be such that $\sigma(y)$ is an end of H. Thus $\sigma(y) = \omega_y$, and ω_y contains the ray Q_y. As before, the definition of Q_y and the fact that $\sigma(a,b) \cap \dot{X} \subseteq D$ imply that $\sigma(y) \in \widehat{D}$. Thus $\sigma(a,b) \subseteq \widehat{D}$, and we have shown that σ is continuous. $\square$

Theorem 5.2 *Every element of the cycle space of an infinite graph G is an edge-disjoint union of cycles in G.*

Proof. Let $Z \in \mathcal{C}(G)$ be given, and let $Z = \sum_{i \in I} Z_i$ where each Z_i is a cycle in G. We first show that I may be partitioned into countable sets I_α so that for all $\alpha \neq \beta$ the graphs $\sum_{i \in I_\alpha} Z_i$ and $\sum_{i \in I_\beta} Z_i$ are edge-disjoint. To do this, consider the auxiliary graph G' with vertex set I in which $i \neq j$ are joined if Z_i and Z_j share an edge. As each Z_i has only countably many edges and each edge lies in only finitely many Z_i, each i has only countably many neighbours in G'. Thus every component of G' is countable, and so the vertex sets I_α of the components of G' form a partition of I with the desired properties. Hence, to prove the theorem, we may assume that I itself is countable. Lemmas 2.8 and 4.4 now imply that there is a countable subgraph H of G such that every Z_i is a cycle in H, and thus Z is an element of the cycle space of H.

Let us rename H as H^0 and Z as Z^0, so that from now on we may use 'H' and 'Z' as variables in Lemma 5.1. Our aim is to write Z^0 as an edge-disjoint union of cycles $C^1, C^2, \dots$ in G. We shall find these C^n inductively inside $Z^{n-1} := \sum_{i \in I} Z_i + C^1 + \dots + C^{n-1}$ by applying Lemma 5.1 to $Z = Z^{n-1}$ in a suitable subgraph H^{n-1} of G. (Thus $C^n \subseteq Z^{n-1}$, and hence $Z^0 \supset Z^1 \supset Z^2 \supset \dots$ with $Z^n = Z^{n-1} + C^n$.)

Starting our inductive definition of the C^n at $n = 1$, let us assume that $C^1, \dots, C^{n-1}$ (and hence $Z^0, \dots, Z^{n-1}$) have been defined as above, and that H^{n-1} is some countable subgraph of G in which $C^1, \dots, C^{n-1}$ and all the Z_i are cycles. To define C^n, let P be as provided by Lemma 5.1 for $H = H^{n-1}$ and $Z = Z^{n-1}$, where $e = vw$ is taken to be the first edge in Z^{n-1} from some fixed enumeration of all the edges of Z^0. (As e will lie in C^n, this choice of e ensures that all the C^n together exhaust Z^0.) The image $\pi_{HG}(P)$ of P in $\overline{G}$ under the canonical projection $\pi_{HG} : \overline{H} \to \overline{G}$ is a path in $\overline{G}$ from v to w. Apply Lemma 3.7 to find an arc $A \subseteq \pi_{HG}(P)$ in $\overline{G}$ with endpoints v and w. Then $A \cup e$ is a circle in $\overline{G}$ whose cycle (in G) is a subgraph of Z^{n-1} containing e, because $P \cap G = P \cap H \subseteq Z^{n-1} - e$; we take C^n to be this cycle.

By Lemma 2.8 there is a countable subgraph H' of G such that C^n is a cycle in H'. By Lemma 4.4 and our assumptions on H^{n-1}, all of $C^1, \dots, C^n$ and all the Z_i then are cycles in $H^n := H^{n-1} \cup H'$, as well as in G.

This completes the inductive definition of the cycles C^n. Since each C^n is a subgraph of Z^{n-1} and $Z^n = Z^{n-1} + C^n$, no edge of C^n is left in Z^n, and so the C^n are indeed edge-disjoint. By the choice of the edges $e = vw$, every edge of $Z = Z^0$ lies in some C^n, and the theorem follows. □

As mentioned before, the cycle space of a graph is not obviously closed under taking infinite sums. Indeed, let $(Z_i)_{i\in I}$ be a thin family of elements of $\mathcal{C}(G)$ (so that $Z := \sum_{i\in I} Z_i$ is well defined), and for each i let $Z_i = \sum_{j\in J_i} C_i^j$ where all the C_i^j are cycles. Then the canonical way to establish Z as an element of $\mathcal{C}(G)$ would be to write it as the 'sum' $Z = \sum_{i\in I,\, j\in J_i} C_i^j$. But this 'sum' may be ill-defined, since the family of all the cycles C_i^j need not be thin even though $(Z_i)_{i\in I}$ is a thin family. For example, if a vertex v lies on exactly two cycles C_i^j for each i, and if both these cycles contain the same two edges at v, then v is not a vertex of Z_i (since we suppress isolated vertices in our definition of sum) and hence does not contradict the thinness of the family $(Z_i)_{i\in I}$; but it does prevent the family of all the C_i^j from being thin.

This phenomenon does not occur, however, when the cycles C_i^j in each of the sums $Z_i = \sum_{j\in J_i} C_i^j$ are edge-disjoint: then $V(Z_i) = \bigcup_{j\in J_i} V(C_i^j)$, and hence if both $(Z_i)_{i\in I}$ and all the $(C_i^j)_{j\in J_i}$ are thin families then so is $(C_i^j)_{i\in I,\, j\in J_i}$. Theorem 5.2 therefore implies that $\mathcal{C}(G)$ is indeed closed under infinite as well as finite sums:

Corollary 5.3 *The cycle space of an infinite graph is closed under taking sums.* □

We now turn to our second result of this section, a cycle-cut orthogonality characterization of the cycle space generalizing Theorem 7.1 of [5] to arbitrary infinite graphs. Recall that a *cut* in G is a set $E(A, B)$ of all the edges of G between the two classes A and B of some bipartition of $V(G)$. A set S of vertices *covers* a cut F if every edge in F has a vertex in S. We say that F is *finitely covered* if there exists a finite set of vertices covering F.

Theorem 5.4 *Let G be any infinite graph, and let $Z \subseteq G$ be any subgraph without isolated vertices. Then the following statements are equivalent:*

(i) $Z \in \mathcal{C}(G)$*;*

(ii) *for every finitely covered cut F of G, $|E(Z) \cap F|$ is (finite and) even.*

Proof. The proof of the implication (i)→(ii) is essentially the same as that of the (i)→(ii) part of Theorem 7.1 in [5]. We now first have to prove that $|E(Z) \cap F|$ is finite, but this is clear since F is covered by finitely many vertices and $Z \in \mathcal{C}(G)$ implies that Z is locally finite.

If the graph G is countable then also the proof of the converse implication is similar to that of [5, Thm. 7.1], except that we now use a normal rather than any end-faithful spanning tree. Every edge $f = tt'$ of a normal spanning tree T of G has the property that every edge of G between the two components of $T - f$ has an endvertex among the finitely many vertices below t and t' in T. Therefore the cut of G associated with f (ie. the set of edges of G between the two components of $T - f$) is finitely covered.

To prove (ii)→(i) for countable G, assume without loss of generality that G is connected and let T be a normal spanning tree in G. Assuming (ii), we show that Z is equal to the sum $Z' \in \mathcal{C}(G)$ of all the fundamental cycles C_e with $e \in E(Z) \setminus E(T)$. For every chord $e \in E(G)$ of T, clearly $e \in Z$ if and only if $e \in Z'$. So consider an edge $f \in T$. Let E_f be the set of edges $e \neq f$ of G between the two components of $T - f$. As shown above, E_f is finitely covered. Since $f \in C_e$ for precisely those chords e of T that lie in E_f, we have $f \in Z'$ if and only if $|E_f \cap E(Z)|$ is odd. By (ii), the latter holds if and only if $f \in Z$, as desired.

The basic idea in our proof of (ii)→(i) for arbitrary graphs G (which need not have normal spanning trees) is to decompose Z into suitable countable subgraphs Y, to extend these to countable separation-faithful subgraphs H by Lemma 2.7, to use the countable case of (ii)→(i) to deduce that $Y \in \mathcal{C}(H) \subseteq \mathcal{C}(G)$ (cf. Corollary 2.6), and finally to combine these results to give $Z \in \mathcal{C}(G)$.

More precisely, let us prove the following claim.

Let X be a subgraph of G without isolated vertices. Suppose that X meets every finitely covered cut of G in an even number of edges. Let X' be any component of X. Then X has a subgraph $Y \in \mathcal{C}(G)$ that is a union of components of X and contains X'. $(*)$

To prove $(*)$, note first that X is locally finite (with even degrees). Hence every component of X is countable. Define a sequence $H_0 \subseteq H_1 \subseteq \dots$ of countable subgraphs of G as follows. Put $H_0 := X'$. Having defined H_i, let H_i' be the graph obtained from H_i by adding every component of X that meets H_i. Then define H_{i+1} as the graph obtained from H_i' by adding, for every finite set $S \subseteq V(H_i)$ and for every pair of distinct components D_1, D_2 of $H_i' - S$ that are contained in a common component D of $G - S$, a D_1–D_2 path in D. Put $H := \bigcup_{i \in \mathbb{N}} H_i$, and let Y be the union of all the components of X that meet H. Then H is countable, $Y \subseteq H$, and H is separation-faithful in G (see the proof of Lemma 2.7 for details).

Let us show that Y satisfies (ii) in H, ie. that Y meets every cut F of H that is covered by a finite set $S \subseteq V(H)$ in an even number of edges. Let such F and S be given, and let $V(H) = A_1 \cup A_2$ be the bipartition of $V(H)$ associated with F. Then every component of $H - S$ has all its vertices in one A_i. Since H is separation-faithful, each component of $G - S$ contains at most one component of $H - S$, and so it meets at most one of the A_i. Let B_1 be the union of A_1 with (the vertex sets of) all the components of $G - S$ avoiding A_2, and let B_2 be the union of A_2 with all the other components of $G - S$. Then $E(B_1, B_2)$ is a cut of G, covered by S. Therefore every edge of $E(B_1, B_2)$ that lies in X must lie in Y (by definition of Y, and as $S \subseteq V(H)$), and hence in H (since $Y \subseteq H$), and hence in F. Thus $|E(Y) \cap F| = |E(X) \cap E(B_1, B_2)|$, and the latter is even by assumption.

Since we have already proved the implication (ii)→(i) for countable graphs we may deduce that Y satisfies (i) in H. Thus, $Y \in \mathcal{C}(H) \subseteq \mathcal{C}(G)$ by Corollary 2.6 (iii), completing the proof of $(*)$.

Let us now use $(*)$ to prove (ii)→(i) for an arbitrary graph G. Let $Z \subseteq G$ be a subgraph that has no isolated vertices and satisfies (ii). Fix a well-ordering of the components of Z. Let us decompose Z into a family of subgraphs $Y_\alpha \in \mathcal{C}(G)$ as in $(*)$, to be defined inductively as follows.

To define Y_0, we apply $(*)$ with $X := Z$ to the first component X' of Z in our well-ordering, and let Y_0 be the graph $Y \in \mathcal{C}(G)$ obtained. Thus, Y_0 satisfies (i) in G. Since we have already shown (i)→(ii) for arbitrary graphs, we may deduce that Y_0 satisfies (ii). Thus Y_0 meets every finitely covered cut of G in an even number of edges, and hence so does $Z - Y_0$.

To define Y_1, we now take X' to be the first component of Z not contained in Y_0, and consider $X := Z - Y_0$. This time, $(*)$ yields a subgraph $Y_1 \in \mathcal{C}(G)$ of $Z - Y_0$. As before, Y_1 and hence also $Z - Y_0 - Y_1$ meets every finitely covered cut of G in an even number of edges.

We continue transfinitely until we have found a sequence (Y_α) of disjoint subgraphs of Z whose union is Z and which all lie in $\mathcal{C}(G)$. Since Z is the sum of all the Y_α, Corollary 5.3 implies that $Z \in \mathcal{C}(G)$. □

6 An open problem

The subgraphs C of a finite graph G that are cycles or other elements of the cycle space of G are easily characterized without any reference to a notion of cyclicity (such as cyclic sequences of vertices etc.). For example, C is a cycle if and only if it is 2-regular and connected, and C is an element of $\mathcal{C}(G)$ if and only if all its vertices have even degree. Similarly, $C \in \mathcal{C}(G)$ if and only if C is orthogonal to every cut of G, ie. meets every cut in an even number of edges.

Since our definition of an infinite cycle appeals to an external notion of cyclicity in an even stronger sense by making reference to the topology of the unit circle, it seems all the more desirable to have similar characterizations for infinite cycles:

Problem *Characterize the cycles and the elements of the cycle space in an infinite graph in purely combinatorial terms.*

Theorem 5.4 offers such a characterization in terms of cuts. Alternatively, one might try to extend the finite 'even degree' characterization of the cycle space to infinite graphs. Clearly, any such characterization will have to refer to ends, but the idea is that such reference should not explicitly appeal to the topology on $\overline{G}$. For example, one might try to define the 'degree' of an end of G in such a way that a subgraph C of G lies in $\mathcal{C}(G)$ if and only if all its vertices have even degree and all its ends have even or infinite degree. One of the problems with such an approach will be in which subgraph to measure the 'degrees' of these ends: probably not in G itself (since an end ω of G that lies on C can contain further rays that have little to do with C), and certainly not in C (where ω will typically split up into many unrelated ends).

References

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

[2] R. Diestel, *Graph Theory* (2nd edition), Springer-Verlag 2000.
`http://www.math.uni-hamburg.de/home/diestel/books/graph.theory/download.html`

[3] R. Diestel, The end structure of a graph: recent results and open problems, *Discrete Mathematics* **100** (1992), 313–327.

[4] R. Diestel, Spanning trees and k-connectedness, *J. Combin. Theory B* **56** (1992), 263–277.

[5] R. Diestel and D. Kühn, On infinite cycles I, *Combinatorica* (to appear).

[6] R. Diestel and D. Kühn, Topological paths, cycles and spanning trees in infinite graphs, submitted.

[7] R. Diestel and D. Kühn, Graph-theoretical versus topological ends of graphs, submitted.

[8] R. Diestel and I. Leader, Normal spanning trees, Aronszajn trees and excluded minors, *J. London Math. Soc.* (2) **63** (2001), 16–32.

[9] R. Diestel and I. Leader, A proof of the bounded graph conjecture, *Invent. math.* **108** (1992), 131–162.

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

Reinhard Diestel and Daniela Kühn, Mathematisches Seminar, Universität Hamburg, Bundesstraße 55, D - 20146 Hamburg, Germany.