Geodetic topological cycles in locally finite graphs

Agelos Georgakopoulos[*]
Mathematisches Seminar
Universität Hamburg
Bundesstraße 55
20146 Hamburg
Germany
georgakopoulos@math.uni-hamburg.de

Philipp Sprüssel
Mathematisches Seminar
Universität Hamburg
Bundesstraße 55
20146 Hamburg
Germany
spruessel@math.uni-hamburg.de

Submitted: Oct 31, 2008; Accepted: Nov 23, 2009; Published: Nov 30, 2009
Mathematics Subject Classification: 05C63

Abstract

We prove that the topological cycle space $\mathcal{C}(G)$ of a locally finite graph G is generated by its geodetic topological circles. We further show that, although the finite cycles of G generate $\mathcal{C}(G)$, its finite geodetic cycles need not generate $\mathcal{C}(G)$.

1 Introduction

A finite cycle C in a graph G is called *geodetic* if, for any two vertices $x, y \in C$, the length of at least one of the two x–y arcs on C equals the distance between x and y in G. It is easy to prove (see Section 3.1):

Proposition 1.1. *The cycle space of a finite graph is generated by its geodetic cycles.*

Our aim is to generalise Proposition 1.1 to the topological cycle space of locally finite infinite graphs.

The *topological cycle space* $\mathcal{C}(G)$ of a locally finite graph G was introduced by Diestel and Kühn [10, 11]. It is built not just from finite cycles, but also from infinite *circles*: homeomorphic images of the unit circle S^1 in the topological space $|G|$ consisting of G, seen as a 1-complex, together with its ends. (See Section 2 for precise definitions.) This space $\mathcal{C}(G)$ has been shown [2, 3, 4, 5, 10, 15] to be the appropriate notion of the cycle space for a locally finite graph: it allows generalisations to locally finite graphs of most of the well-known theorems about the cycle space of finite graphs, theorems which fail for

[*]Supported by GIF grant no. I-879-124.6.

infinite graphs if the usual finitary notion of the cycle space is applied. It thus seems that the topological cycle space is an important object that merits further investigation. (See [6, 7] for introductions to the subject.)

As in the finite case, one fundamental question is which natural subsets of the topological cycle space generate it, and how. It has been shown, for example, that the fundamental circuits of topological spanning trees do (but not those of arbitrary spanning trees) [10], or the non-separating induced cycles [2], or that every element of $\mathcal{C}(G)$ is a sum of disjoint circuits [11, 7, 19]—a trivial observation in the finite case, which becomes rather more difficult for infinite G. (A shorter proof, though still non-trivial, is given in [15].) Another standard generating set for the cycle space of a finite graph is the set of geodetic cycles (Proposition 1.1), and it is natural to ask whether these still generate $\mathcal{C}(G)$ when G is infinite.

But what is a geodetic topological circle? One way to define it would be to apply the standard definition, stated above before Proposition 1.1, to arbitrary circles, taking as the length of an arc the number of its edges (which may now be infinite). As we shall see, Proposition 1.1 will fail with this definition, even for locally finite graphs. Indeed, with hindsight we can see why it should fail: when G is infinite then giving every edge length 1 will result in path lengths that distort rather than reflect the natural geometry of $|G|$: edges 'closer to' ends must be shorter, if only to give paths between ends finite lengths.

It looks, then, as though the question of whether or not Proposition 1.1 generalises might depend on how exactly we choose the edge lengths in our graph. However, our main result is that this is not the case: we shall prove that no matter how we choose the edge lengths, as long as the resulting arc lengths induce a metric compatible with the topology of $|G|$, the geodetic circles in $|G|$ will generate $\mathcal{C}(G)$. Note, however, that the question of which circles are geodetic does depend on our choice of edge lengths, even under the assumption that a metric compatible with the topology of $|G|$ is induced.

If $\ell : E(G) \to \mathbb{R}_+$ is an assignment of edge lengths that has the above property, we call the pair $(|G|, \ell)$ a *metric representation of* G. We then call a circle C *ℓ-geodetic* if for any points x, y on C the distance between x and y in C is the same as the distance between x and y in $|G|$. See Section 2.2 for precise definitions and more details.

We can now state the main result of this paper more formally:

Theorem 1.2. *For every metric representation $(|G|, \ell)$ of a connected locally finite graph G, the topological cycle space $\mathcal{C}(G)$ of G is generated by the ℓ-geodetic circles in G.*

Motivated by the current work, the first author initiated a more systematic study of topologies on graphs that can be induced by assigning lengths to the edges of the graph. In this context, it is conjectured that Theorem 1.2 generalises to arbitrary compact metric spaces if the notion of the topological cycle space is replaced by an analogous homology [14].

We prove Theorem 1.2 in Section 4, after giving the required definitions and basic facts in Section 2 and showing that Proposition 1.1 holds for finite graphs but not for infinite ones in Section 3. Finally, in Section 5 we will discuss some further problems.

2 Definitions and background

2.1 The topological space $|G|$ and $\mathcal{C}(G)$

Unless otherwise stated, we will be using the terminology of [7] for graph-theoretical concepts and that of [1] for topological ones. Let $G = (V, E)$ be a *locally finite* graph — i.e. every vertex has a finite degree — finite or infinite, fixed throughout this section.

The graph-theoretical distance between two vertices $x, y \in V$, is the minimum $n \in \mathbb{N}$ such that there is an x–y path in G comprising n edges. Unlike the frequently used convention, we will not use the notation $d(x, y)$ to denote the graph-theoretical distance, as we use it to denote the distance with respect to a metric d on $|G|$.

A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*. A *tail of* a ray R is an infinite 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 $\Omega = \Omega(G)$, and we define $\hat{V} := V \cup \Omega$.

Let G bear the topology of a 1-complex, where the 1-cells are real intervals of arbitrary lengths[1]. 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 $C = C(S, \omega)$ 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 $C(S, \omega)$. 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 [9]) that in fact $|G|$ is the Freudenthal compactification [13] of the 1-complex G.

A continuous map σ from the real unit interval $[0, 1]$ to $|G|$ is a *topological path* in $|G|$; the images under σ of 0 and 1 are its *endpoints*. A homeomorphic image of the real unit interval in $|G|$ is an *arc* in $|G|$. Any set $\{x\}$ with $x \in |G|$ is also called an arc in $|G|$. A homeomorphic image of S^1, the unit circle in $\mathbb{R}^2$, in $|G|$ is a (*topological cycle* or) *circle* in $|G|$. Note that any arc, circle, cycle, path, or image of a topological path is closed in $|G|$, since it is a continuous image of a compact space in a Hausdorff space.

A subset D of E is a *circuit* if there is a circle C in $|G|$ such that $D = \{e \in E \mid e \subseteq C\}$. Call a family $\mathcal{F} = (D_i)_{i \in I}$ of subsets of E *thin* if no edge lies in D_i for infinitely many indices i. Let the *(thin) sum* $\sum \mathcal{F}$ of this family be the set of all edges that lie in D_i for an odd number of indices i, and let the *topological cycle space* $\mathcal{C}(G)$ of G be the set of all sums of thin families of circuits. In order to keep our expressions simple, we will, with a slight abuse, not stricly distinguish circles, paths and arcs from their edge sets.

[1]Every edge is homeomorphic to a real closed bounded interval, 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 the topology does not depend on the lengths of the intervals homeomorphic to edges.

2.2 Metric representations

Suppose that the lengths of the 1-cells (edges) of the locally finite graph G are given by a function $\ell : E(G) \to \mathbb{R}_+$. Every arc in $|G|$ is either a subinterval of an edge or the closure of a disjoint union of open edges or half-edges (at most two, one at either end), and we define its *length* as the length of this subinterval or as the (finite or infinite) sum of the lengths of these edges and half-edges, respectively. Given two points $x, y \in |G|$, write $d_\ell(x, y)$ for the infimum of the lengths of all x–y arcs in $|G|$. It is straightforward to prove:

Proposition 2.1. *If for every two points $x, y \in |G|$ there is an x-y arc of finite length, then d_ℓ is a metric on $|G|$.*

This metric d_ℓ will in general not induce the topology of $|G|$. If it does, we call $(|G|, \ell)$ a *metric representation* of G (other topologies on a graph that can be induced by edge lengths in a similar way are studied in [14]). We then call a circle C in $|G|$ *ℓ-geodetic* if, for every two points $x, y \in C$, one of the two x–y arcs in C has length $d_\ell(x, y)$. If C is ℓ-geodetic, then we also call its circuit *ℓ-geodetic*.

Metric representations do exist for every locally finite graph G. Indeed, pick a normal spanning tree T of G with root $x \in V(G)$ (its existence is proved in [7, Theorem 8.2.4]), and define the length $\ell(uv)$ of any edge $uv \in E(G)$ as follows. If $uv \in E(T)$ and $v \in xTu$, let $\ell(uv) = 1/2^{|xTu|}$. If $uv \notin E(T)$, let $\ell(uv) = \sum_{e \in uTv} \ell(e)$. It is easy to check that d_ℓ is a metric of $|G|$ inducing its topology [8].

2.3 Basic facts

In this section we give some basic properties of $|G|$ and $\mathcal{C}(G)$ that we will need later.

One of the most fundamental properties of $\mathcal{C}(G)$ is that:

Lemma 2.2 ([11])**.** *For any locally finite graph G, every element of $\mathcal{C}(G)$ is an edge-disjoint sum of circuits.*

As already mentioned, $|G|$ is a compactification of the 1-complex G:

Lemma 2.3 ([7, Proposition 8.5.1])**.** *If G is locally finite and connected, then $|G|$ is a compact Hausdorff space.*

The next statement follows at once from Lemma 2.3.

Corollary 2.4. *If G is locally finite and connected, then the closure in $|G|$ of an infinite set of vertices contains an end.*

The following basic fact can be found in [16, p. 208].

Lemma 2.5. *The image of a topological path with endpoints x, y in a Hausdorff space X contains an arc in X between x and y.*

As a consequence, being linked by an arc is an equivalence relation on $|G|$; a set $Y \subset |G|$ is called *arc-connected* if Y contains an arc between any two points in Y. Every arc-connected subspace of $|G|$ is connected. Conversely, we have:

Lemma 2.6 ([12])**.** *If G is a locally finite graph, then every closed connected subspace of $|G|$ is arc-connected.*

The following lemma is a standard tool in infinite graph theory.

Lemma 2.7 (König's Infinity Lemma [17])**.** *Let $V_0, V_1, \ldots$ be an infinite sequence of disjoint non-empty finite sets, and let G be a graph on their union. Assume that every vertex v in a set V_n with $n \geqslant 1$ has a neighbour in V_{n-1}. Then G contains a ray $v_0 v_1 \cdots$ with $v_n \in V_n$ for all n.*

3 Generating $\mathcal{C}(G)$ by geodetic cycles

3.1 Finite graphs

In this section finite graphs, like infinite ones, are considered as 1-complexes where the 1-cells (i.e. the edges) are real intervals of arbitrary lengths. Given a metric representation $(|G|, \ell)$ of a finite graph G, we can thus define the *length* $\ell(X)$ of a path or cycle X in G by $\ell(X) = \sum_{e \in E(X)} \ell(e)$. Note that, for finite graphs, any assignment of edge lengths yields a metric representation. A cycle C in G is *ℓ-geodetic*, if for any $x, y \in V(C)$ there is no x–y path in G of length strictly less than that of each of the two x–y paths on C.

The following theorem generalises Proposition 1.1.

Theorem 3.1. *For every finite graph G and every metric representation $(|G|, \ell)$ of G, every cycle C of G can be written as a sum of ℓ-geodetic cycles of length at most $\ell(C)$.*

Proof. Suppose that the assertion is false for some $(|G|, \ell)$, and let D be a cycle in G of minimal length among all cycles C that cannot be written as a sum of ℓ-geodetic cycles of length at most $\ell(C)$. As D is not ℓ-geodetic, it is easy to see that there is a path P with both endvertices on D but no inner vertex in D that is shorter than the paths Q_1, Q_2 on D between the endvertices of P. Thus D is the sum of the cycles $D_1 := P \cup Q_1$ and $D_2 := P \cup Q_2$. As D_1 and D_2 are shorter than D, they are each a sum of ℓ-geodetic cycles of length less than $\ell(D)$, which implies that D itself is such a sum, a contradiction. □

By letting all edges have length 1, Theorem 3.1 implies Proposition 1.1.

3.2 Failure in infinite graphs

As already mentioned, Proposition 1.1 does not naively generalise to locally finite graphs: there are locally finite graphs whose topological cycle space contains a circuit that is not a thin sum of circuits that are geodetic in the traditional sense, i.e. when every edge has length 1. Such a counterexample is given in Figure 3.1. The graph H shown there is a subdivision of the *infinite ladder*; the infinite ladder is a union of two rays $R_x = x_1x_2\cdots$ and $R_y = y_1y_2\cdots$ plus an edge x_ny_n for every $n \in \mathbb{N}$, called the *n-th rung* of the ladder. By subdividing, for every $n \geqslant 2$, the n-th rung into $2n$ edges, we obtain H. For every $n \in \mathbb{N}$, the (unique) shortest x_n–y_n path contains the first rung e and has length $2n-1$. As every circle (finite or infinite) must contain the subdivision of at least one rung, every geodetic circuit contains e. On the other hand, Figure 3.1 shows an element C of $\mathcal{C}(H)$ that contains infinitely many rungs. As every circle can contain at most two rungs, we need an infinite family of geodetic circuits to generate C, but since they all have to contain e the family cannot be thin.

The graph H is however not a counterexample to Theorem 1.2, since the constant edge lengths 1 do not induce a metric of $|H|$.

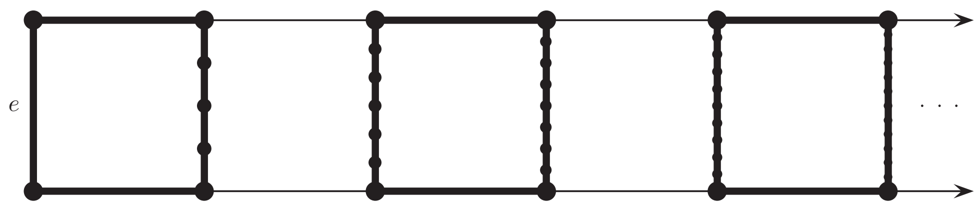

Fig. 3.1: A 1-ended graph and an element of its topological cycle space (drawn thick) which is not the sum of a thin family of geodetic circuits.

4 Generating $\mathcal{C}(G)$ by geodetic circles

Let G be an arbitrary connected locally finite graph, finite or infinite, consider a fixed metric representation $(|G|, \ell)$ of G and write $d = d_\ell$. We want to assign a length to every arc or circle, but also to other objects like elements of $\mathcal{C}(G)$. To this end, let X be an arc or circle in $|G|$, an element of $\mathcal{C}(G)$, or the image of a topological path in $|G|$. It is easy to see that for every edge e, $e \cap X$ is the union of at most two subintervals of e and thus has a natural length which we denote by $\ell(e \cap X)$; moreover, X is the closure in $|G|$ of $\bigcup_{e \in G}(e \cap X)$ (unless X contains less than two points). We can thus define the *length* of X as $\ell(X) := \sum_{e \in G} \ell(e \cap X)$.

Note that not every such X has finite length (see Section 5). But the length of an ℓ-geodetic circle C is always finite. Indeed, as $|G|$ is compact, there is an upper bound ε_0 such that $d(x, y) \leqslant \varepsilon_0$ for all $x, y \in |G|$. Therefore, C has length at most $2\varepsilon_0$.

For the proof of Theorem 1.2 it does not suffice to prove that every circuit is a sum of a thin family of ℓ-geodetic circuits. (Moreover, the proof of the latter statement turns out to be as hard as the proof of Theorem 1.2.) For although every element C of $\mathcal{C}(G)$ is a sum of a thin family of circuits (even of finite circuits, see [7, Corollary 8.5.9]), representations of all the circles in this family as sums of thin families of ℓ-geodetic circuits will not necessarily combine to a similar representation for C, because the union of infinitely many thin families need not be thin.

In order to prove Theorem 1.2, we will use a sequence $\hat{S}_i$ of finite auxiliary graphs whose limit is G. Given an element C of $\mathcal{C}(G)$ that we want to represent as a sum of ℓ-geodetic circuits, we will for each i consider an element $C|\hat{S}_i$ of the cycle space of $\hat{S}_i$ induced by C — in a way that will be made precise below — and find a representation of $C|\hat{S}_i$ as a sum of geodetic cycles of $\hat{S}_i$, provided by Theorem 3.1. We will then use the resulting sequence of representations and compactness to obtain a representation of C as a sum of ℓ-geodetic circuits.

4.1 Restricting paths and circles

To define the auxiliary graphs mentioned above, pick a vertex $w \in G$, and let, for every $i \in \mathbb{N}$, S_i be the set of vertices of G whose graph-theoretical distance from w is at most i; also let $S_{-1} = \emptyset$. Note that $S_0 = \{w\}$, every S_i is finite, and $\bigcup_{i \in \mathbb{N}} S_i = V(G)$. For every $i \in \mathbb{N}$, define $\tilde{S}_i$ to be the subgraph of G on S_{i+1}, containing those edges of G that are incident with a vertex in S_i. Let $\hat{S}_i$ be the graph obtained from $\tilde{S}_i$ by joining every two vertices in $S_{i+1} - S_i$ that lie in the same component of $G - S_i$ with an edge; these new edges are the *outer edges* of $\hat{S}_i$. For every $i \in \mathbb{N}$, a metric representation $(|\hat{S}_i|, \ell_i)$ can be defined as follows: let every edge e of $\hat{S}_i$ that also lies in $\tilde{S}_i$ have the same length as in $|G|$, and let every outer edge $e = uv$ of $\hat{S}_i$ have length $d_\ell(u, v)$. For any two points $x, y \in |\hat{S}_i|$ we will write $d_i(x, y)$ for $d_{\ell_i}(x, y)$ (the latter was defined at the end of Section 2.1). Recall that in the previous subsection we defined a length $\ell_i(X)$ for every path, cycle, element of the cycle space, or image of a topological path X in $|\hat{S}_i|$.

If X is an arc with endpoints in $\hat{V}$ or a circle in $|G|$, define the *restriction* $X|\hat{S}_i$ *of* X *to* $\hat{S}_i$ as follows. If X avoids S_i, let $X|\hat{S}_i = \emptyset$. Otherwise, start with $E(X) \cap E(\hat{S}_i)$ and add all outer edges uv of $\hat{S}_i$ such that X contains a u–v arc that meets $\hat{S}_i$ only in u and v. We defined $X|\hat{S}_i$ to be an edge set, but we will, with a slight abuse, also use the same term to denote the subgraph of $\hat{S}_i$ spanned by this edge set. Clearly, the restriction of a circle is a cycle and the restriction of an arc is a path. For a path or cycles X in $\hat{S}_j$ with $j > i$, we define the restriction $X|\hat{S}_i$ to $\hat{S}_i$ analogously.

Note that in order to obtain $X|\hat{S}_i$ from X, we deleted a set of edge-disjoint arcs or paths in X, and for each element of this set we put in $X|\hat{S}_i$ an outer edge with the same endpoints. As no arc or path is shorter than an outer edge with the same endpoints, we easily obtain:

Lemma 4.1. *Let $i \in \mathbb{N}$ and let X be an arc or a circle in $|G|$ (respectively, a path or cycle in $\hat{S}_j$ with $j > i$). Then $\ell_i(X|\hat{S}_i) \leqslant \ell(X)$ (resp. $\ell_i(X|\hat{S}_i) \leqslant \ell_j(X)$).*

A consequence of this is the following:

Lemma 4.2. *If $x, y \in S_{i+1}$ and P is a shortest x–y path in $\hat{S}_i$ with respect to ℓ_i then $\ell_i(P) = d(x, y)$.*

Proof. Suppose first that $\ell_i(P) < d(x, y)$. Replacing every outer edge uv in P by a u–v arc of length $\ell_i(uv) + \varepsilon$ in $|G|$ for a sufficiently small ε, we obtain a topological x–y path in $|G|$ whose image is shorter than $d(x, y)$. Since, by Lemma 2.5, the image of every topological

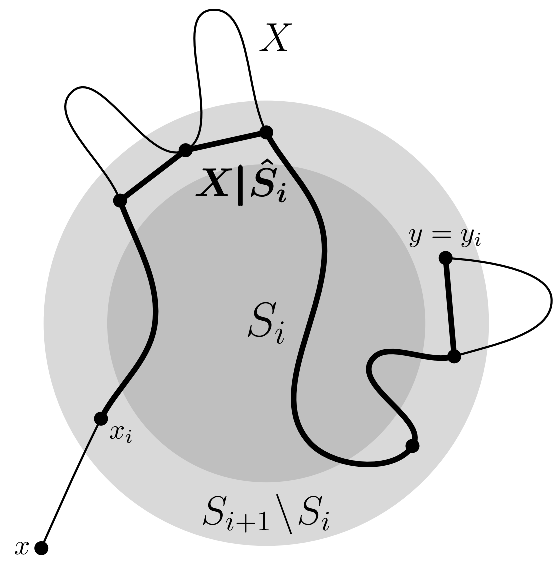

Fig. 4.2: The restriction of an x–y arc X to the x_i–y_i path $X|\hat{S}_i$.

path contains an arc with the same endpoints, this contradicts the definition of $d(x, y)$. Next, suppose that $\ell_i(P) > d(x, y)$. In this case, there is by the definition of $d(x, y)$ an x–y arc Q in $|G|$ with $\ell(Q) < \ell_i(P)$. Then $\ell_i(Q|\hat{S}_i) \leqslant \ell(Q) < \ell_i(P)$ by Lemma 4.1, contradicting the choice of P. $\square$

Let $C \in \mathcal{C}(G)$. For the proof of Theorem 1.2 we will construct a family of ℓ-geodetic circles in ω steps, choosing finitely many of these at each step. To ensure that the resulting family will be thin, we will restrict the lengths of those circles: the next two lemmas will help us bound these lengths from above, using the following amounts ε_i that vanish as i grows.

$$\varepsilon_i := \sup\{d(x, y) \mid x, y \in |G| \textit{ and there is an } x\text{–}y \textit{ arc in } |G| \setminus G[S_{i-1}]\}.$$

The space $|G| \setminus G[S_{i-1}]$ considered in this definition is the same as the union of $|G - S_{i-1}|$ and the inner points of all edges from S_{i-1} to $V(G) \setminus S_{i-1}$. Note that as $|G|$ is compact, each ε_i is finite.

Lemma 4.3. *Let $j \in \mathbb{N}$, let C be a cycle in $\hat{S}_j$, and let $i \in \mathbb{N}$ be the smallest index such that C meets S_i. Then C can be written as a sum of ℓ_j-geodetic cycles in $\hat{S}_j$ each of which has length at most $5\varepsilon_i$ in $\hat{S}_j$.*

Proof. We will say that a cycle D in $\hat{S}_j$ is *a C-sector* if there are vertices x, y on D such that one of the x–y paths on D has length at most ε_i and the other, called a *C-part* of D, is contained in C.

We claim that every C-sector D longer than $5\varepsilon_i$ can be written as a sum of cycles shorter than D, so that every cycle in this sum either has length at most $5\varepsilon_i$ or is another C-sector. Indeed, let Q be a C-part of D and let x, y be its endvertices. Every edge e of

Q has length at most $2\varepsilon_i$: otherwise the midpoint of e has distance greater than ε_i from each endvertex of e, contradicting the definition of ε_i. As Q is longer than $4\varepsilon_i$, there is a vertex z on Q whose distance, with respect to ℓ_j, along Q from x is larger than ε_i but at most $3\varepsilon_i$. Then the distance of z from y along Q is also larger than ε_i. By the definition of ε_i and Lemma 4.2, there is a z–y path P in $\hat{S}_j$ with $\ell_j(P) \leqslant \varepsilon_i$.

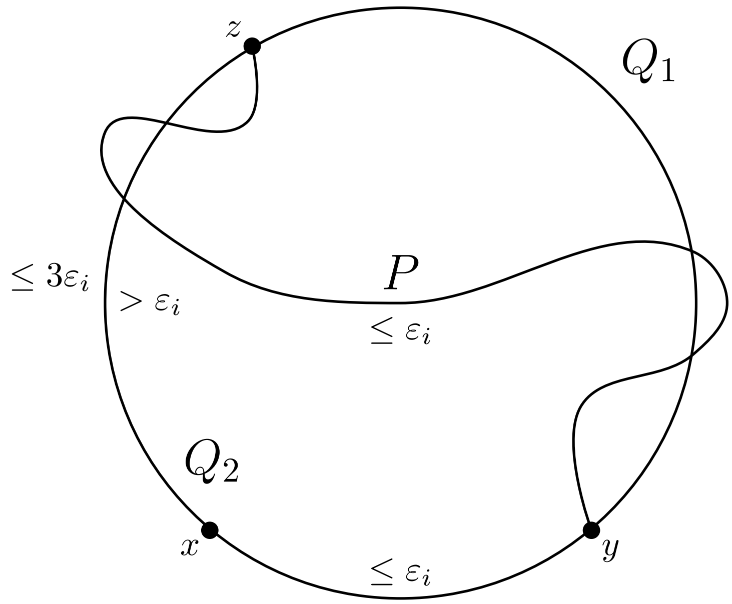

Fig. 4.3: The paths Q_1, Q_2, and P in the proof of Lemma 4.3.

Let $Q_1 = zQy$ and let Q_2 be the other z–y path in D. (See also Figure 4.3.) Note that Q_2 is the concatenation of zQ_2x and xQ_2y. Since $\varepsilon_i < \ell_j(zQ_2x) \leqslant 3\varepsilon_i$ and $\ell_j(xQ_2y) \leqslant \varepsilon_i$, we have $\varepsilon_i < \ell_j(Q_2) \leqslant 4\varepsilon_i$. For any two paths R, L, we write $R+L$ as a shorthand for the symmetric difference of $E(R)$ and $E(L)$. It is easy to check that every vertex is incident with an even number of edges in $Q_2 + P$, which means that $Q_2 + P$ is an element of the cycle space of $\hat{S}_j$, so by Lemma 2.2 it can be written as a sum of edge-disjoint cycles in $\hat{S}_j$. Since $\ell_j(Q_2 + P) \leqslant \ell_j(Q_2) + \ell_j(P) \leqslant 4\varepsilon_i + \varepsilon_i = 5\varepsilon_i$, every such cycle has length at most $5\varepsilon_i$. On the other hand, we claim that $Q_1 + P$ can be written as a sum of C-sectors that are contained in $Q_1 \cup P$. If this is true then each of those C-sectors will be shorter than D since

$$\ell_j(Q_1 \cup P) \leqslant \ell_j(Q_1) + \ell_j(P) \leqslant \ell_j(Q_1) + \varepsilon_i < \ell_j(Q_1) + \ell_j(Q_2) = \ell_j(D).$$

To prove that $Q_1 + P$ is a sum of such C-sectors, consider the vertices in $X := V(Q_1) \cap V(P)$ in the order they appear on P (recall that P starts at z and ends at y) and let v be the last vertex in this order such that $Q_1v + Pv$ is the (possibly trivial) sum of C-sectors contained in $Q_1 \cup P$ (there is such a vertex since $z \in X$ and $Q_1z + Pz = \emptyset$). Suppose $v \neq y$ and let w be the successor of v in X. The paths vQ_1w and vPw have no vertices in common other than v and w, hence either they are edge-disjoint or they both consist of the same edge vw. In both cases, $Q_1w + Pw = (Q_1v + Pv) + (vQ_1w + vPw)$ is the sum of C-sectors contained in $Q_1 \cup P$, since $Q_1 v + Pv$ is such a sum and $vQ_1w + vPw$ is either the empty edge-set or a C-sector contained in $Q_1 \cup P$ (recall that $vQ_1w \subset C$ and $\ell_j(vPw) \leqslant \varepsilon_i$). This contradicts the choice of v, therefore $v = y$ and $Q_1 + P$ is a sum of C-sectors as required.

Thus every C-sector longer than $5\varepsilon_i$ is a sum of shorter cycles, either C-sectors or cycles shorter than $5\varepsilon_i$. As $\hat{S}_j$ is finite and C is a C-sector itself, repeated application of this fact yields that C is a sum of cycles not longer than $5\varepsilon_i$. By Proposition 3.1, every cycle in this sum is a sum of ℓ_j-geodetic cycles in $\hat{S}_j$ not longer than $5\varepsilon_i$; this completes the proof. $\square$

Lemma 4.4. *The sequence $(\varepsilon_i)_{i\in\mathbb{N}}$ converges to zero.*

Proof. The sequence $(\varepsilon_i)_{i\in\mathbb{N}}$ converges since it is decreasing. Suppose there is an $\varepsilon > 0$ with $\varepsilon_i > \varepsilon$ for all i. Thus, for every $i \in \mathbb{N}$, there is a component C_i of $|G| \setminus G[S_i]$ in which there are two points of distance at least ε. For every $i \in \mathbb{N}$, pick a vertex $c_i \in C_i$. By Corollary 2.4, there is an end ω in the closure of the set $\{c_0, c_1, \dots\}$ in $|G|$. Let $\hat{C}(S_i, \omega)$ denote the component of $|G| \setminus G[S_i]$ that contains ω. It is easy to see that the sets $\hat{C}(S_i, \omega)$, $i \in \mathbb{N}$, form a neigbourhood basis of ω in $|G|$.

As $U := \{x \in |G| \mid d(x, \omega) < \frac{1}{2}\varepsilon\}$ is open in $|G|$, it has to contain $\hat{C}(S_i, \omega)$ for some i. Furthermore, there is a vertex $c_j \in \hat{C}(S_i, \omega)$ with $j \geqslant i$, because ω lies in the closure of $\{c_0, c_1, \dots\}$. As $S_j \supset S_i$, the component C_j of $|G| \setminus G[S_j]$ is contained in $\hat{C}(S_i, \omega)$ and thus in U. But any two points in U have distance less than ε, contradicting the choice of C_j. $\square$

This implies in particular that:

Corollary 4.5. *Let $\varepsilon > 0$ be given. There is an $n \in \mathbb{N}$ such that for every $i \geqslant n$, every outer edge of $\hat{S}_i$ is shorter than ε.*

4.2 Limits of paths and cycles

In this section we develop some tools that will help us obtain ℓ-geodetic circles as limits of sequences of ℓ_i-geodetic cycles in the $\hat{S}_i$.

A *chain* of paths (respectively cycles) is a sequence $X_j, X_{j+1}, \dots$ of paths (resp. cycles), such that every X_i with $i \geqslant j$ is the restriction of X_{i+1} to $\hat{S}_i$.

Definition 4.6. *The* limit *of a chain $X_j, X_{j+1}, \dots$ of paths or cycles, is the closure in $|G|$ of the set*

$$\tilde{X} := \bigcup_{j \leqslant i < \omega} \left(X_i \cap \tilde{S}_i \right).$$

Unfortunately, the limit of a chain of cycles does not have to be a circle, as shown in Figure 4.4. However, we are able to prove the following lemma.

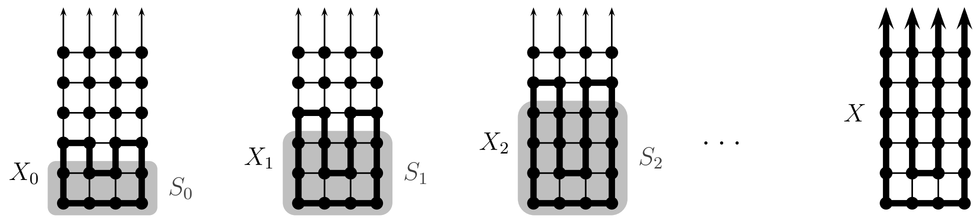

Fig. 4.4: A chain $X_0, X_1, \dots$ of cycles (drawn thick), whose limit X is not a circle (but the edge-disjoint union of two circles).

Lemma 4.7. *The limit of a chain of cycles is a continuous image of S^1 in $|G|$. The limit of a chain of paths is the image of a topological path in $|G|$. The corresponding continuous map can be chosen so that every point in G has at most one preimage, while the preimage of each end of G is a totally disconnected set.*

Proof. Let $X_0, X_1, \dots$ be a chain of cycles (proceed analogously for a chain $X_j, X_{j+1}, \dots$) and let X be its limit. We define the desired map $\sigma : S^1 \to X$ with the help of homeomorphisms $\sigma_i : S^1 \to X_i$ for every $i \in \mathbb{N}$. Start with some homeomorphism $\sigma_0 : S^1 \to X_0$. Now let $i \geqslant 1$ and suppose that $\sigma_{i-1} : S^1 \to X_{i-1}$ has already been defined. We change σ_{i-1} to σ_i by mapping the preimage of any outer edge in X_{i-1} to the corresponding path in X_i. While we do this, we make sure that the preimage of every outer edge in X_i is not longer than $\frac{1}{i}$.

Now for every $x \in S^1$, define $\sigma(x)$ as follows. If there is an $n \in \mathbb{N}$ such that $\sigma_i(x) = \sigma_n(x)$ for every $i \geqslant n$, then define $\sigma(x) = \sigma_n(x)$. Otherwise, $\sigma_i(x)$ lies on an outer edge $u_i v_i$ for every $i \in \mathbb{N}$. By construction, there is exactly one end ω in the closure of $\{u_0, v_0, u_1, v_1, \dots\}$ in $|G|$, and we put $\sigma(x) := \omega$.

It is straightforward to check that $\sigma : S^1 \to X$ is continuous, and that $\tilde{X} \subseteq \sigma(S^1)$. As $\sigma(S^1)$ is a continuous image of the compact space S^1 in the Hausdorff space $|G|$, it is closed in $|G|$, thus $\sigma(S^1) = X$. By construction, only ends can have more than one preimage under σ. Moreover, as we defined σ_i so that the preimage of every outer edge is not longer than $\frac{1}{i}$, any two points $x, y \in S^1$ that are mapped by σ to ends are mapped by σ_i to distinct outer edges for every sufficiently large i. Therefore, there are points in S^1 between x and y that are mapped to a vertex by σ_i and hence also by σ, which shows that the preimage of each end under σ is totally disconnected.

For a chain $X_0, X_1, \dots$ of paths, the construction is slightly different: As the endpoints of the paths X_i may change while i increases, we let $\sigma_i : [0,1] \to X$ map a short interval $[0, \delta_i]$ to the first vertex of X_i, and the interval $[1-\delta_i, 1]$ to the last vertex of X_i, where δ_i is a sequence of real numbers converging to zero. Except for this difference, the construction of a continuous map $\sigma : [0,1] \to X$ imitates that of the previous case. □

Recall that a circle is ℓ-geodetic if for every two points $x, y \in C$, one of the two x–y arcs in C has length $d(x, y)$. Equivalently, a circle C is ℓ-geodetic if it has no *shortcut*, that is, an arc in $|G|$ with endpoints $x, y \in C$ and length less than both x–y arcs in C. Indeed, if a circle C has no shortcut, it is easily seen to contain a shortest x–y arc for any points $x, y \in C$.

It may seem more natural if a shortcut of C is a C-*arc*, that is, an arc that meets C only at its endpoints. The following lemma will allow us to only consider such shortcuts (in particular, we only have to consider arcs with endpoints in $\hat{V}$).

Lemma 4.8. *Every shortcut of a circle C in $|G|$ contains a C-arc which is also a shortcut of C.*

Proof. Let P be a shortcut of C with endpoints x, y. As C is closed, every point in $P \setminus C$ is contained in a C-arc in P. Suppose no C-arc in P is a shortcut of C. We can find a family $(W_i)_{i\in\mathbb{N}}$ of countably many internally disjoint arcs in P, such that for every i, W_i is either a C-arc or an arc contained in C, and every edge in P lies in some W_i (there may, however, exist ends in P that are not contained in any arc W_i). For every i, let x_i, y_i be the endpoints of W_i and pick a x_i–y_i arc K_i as follows. If W_i is contained in C, let $K_i = W_i$. Otherwise, W_i is a C-arc and we let K_i be the shortest x_i–y_i arc on C. Note that since W_i is not a shortcut of C, K_i is at most as long as W_i.

Let K be the union of all the arcs K_i. Clearly, the closure $\overline{K}$ of K in $|G|$ is contained in C, contains x and y, and is at most as long as P. It is easy to see that $\overline{K}$ is a connected topological space; indeed, if not, then there are distinct edges e, f on C, so that both components of $C - \{e, f\}$ meet $\overline{K}$, which cannot be the case by the construction of K. By Theorem 2.6, $\overline{K}$ is also arc-connected, and so it contains an x–y arc that is at most as long as P, contradicting the fact that P is a shortcut of C.

Thus, P contains a C-arc which is also a shortcut of C. □

By the following lemma, the restriction of any geodetic circle is also geodetic.

Lemma 4.9. *Let $i > j$ and let C be an ℓ-geodetic circle in $|G|$ (respectively, an ℓ_i-geodetic cycle in $\hat{S}_i$). Then $C_j := C|\hat{S}_j$ is an ℓ_j-geodetic cycle in $\hat{S}_j$, unless $C_j = \emptyset$.*

Proof. Suppose for contradiction, that C_j has a shortcut P between the vertices x, y. Clearly, x, y lie in C, so let Q_1, Q_2 be the two x–y arcs (resp. x–y paths) in C. We claim that $\ell(Q_k) > d(x, y)$ (resp. $\ell_i(Q_k) > d(x, y)$) for $k = 1, 2$. Indeed, as P is a shortcut of C_j, and $Q_k|\hat{S}_j$ is a subpath of C_j with endvertices x, y for $k = 1, 2$, we have $\ell_j(Q_k|\hat{S}_j) > \ell_j(P)$. Moreover, by Lemma 4.1 we have $\ell(Q_k) \geqslant \ell_j(Q_k|\hat{S}_j)$ (resp. $\ell_i(Q_k) \geqslant \ell_j(Q_k|\hat{S}_j)$), and by Lemma 4.2 $\ell_j(P) \geqslant d(x, y)$, so our claim is proved. But then, by the definition of $d(x, y)$ (resp. by Lemma 4.2), there is an x–y arc Q in $|G|$ such that $\ell(Q) < \ell(Q_k)$ (resp. an x–y path Q in $\hat{S}_i$ such that $d(x, y) = \ell_i(Q) < \ell_i(Q_k)$) for $k = 1, 2$, contradicting the fact that C is ℓ-geodetic (resp. ℓ_i-geodetic). □

As already mentioned, the limit of a chain of cycles does not have to be a circle. Fortunately, the limit of a chain of geodetic cycles is always a circle, and in fact an ℓ-geodetic one:

Lemma 4.10. *Let C be the limit of a chain $C_0, C_1, \ldots$ of cycles, such that, for every $i \in \mathbb{N}$, C_i is ℓ_i-geodetic in $\hat{S}_i$. Then C is an ℓ-geodetic circle.*

Proof. Let σ be the map provided by Lemma 4.7 (with C_i instead of X_i). We claim that σ is injective.

Indeed, as only ends can have more than one preimage under σ, suppose, for contradiction, that ω is an end with two preimages. These preimages subdivide S^1 into two components P_1, P_2. Choose $\varepsilon \in \mathbb{R}_+$ smaller than the lengths of $\sigma(P_1)$ and $\sigma(P_2)$. (Note that as the preimage of ω is totally disconnected, both $\sigma(P_1)$ and $\sigma(P_2)$ have to contain edges of G, so their lengths are positive.) By Corollary 4.5, there is a j such that every outer edge of $\hat{S}_j$ is shorter than ε. On the other hand, for a sufficiently large $i \geqslant j$, the restrictions of $\sigma(P_1)$ and $\sigma(P_2)$ to $\hat{S}_i$ are also longer than ε. Thus, the distance along C_i between the first and the last vertex of $\sigma(P_1)|\hat{S}_i$ is larger than ε. As those vertices lie in the same component of $G - S_i$ (namely, in $C(S_i, \omega)$), there is an outer edge of $\hat{S}_i$ between them. This edge is shorter than ε and thus a shortcut of C_i, contradicting the fact that C_i is ℓ_i-geodetic.

Thus, σ is injective. As any bijective, continuous map between a compact space and a Hausdorff space is a homeomorphism, C is a circle.

Suppose, for contradiction, there is a shortcut P of C between points $x, y \in C \cap \hat{V}$. Choose $\varepsilon > 0$ such that P is shorter by at least 3ε than both x–y arcs on C. Then, there is an i such that the restrictions Q_1, Q_2 of the x–y arcs on C to $\hat{S}_i$ are longer by at least 2ε than $P_i := P|\hat{S}_i$ (Q_1, Q_2 lie in C_i by the definition of σ, but note that they may have different endpoints). By Corollary 4.5, we may again assume that every outer edge of $\hat{S}_i$ is shorter than ε. If x does not lie in $\hat{S}_i$, then the first vertices of P_i and Q_1 lie in the component of $G - S_i$ that contains x (or one of its rays if x is an end). The same is true for y and the last vertices of P_i and Q_1. Thus, we may extend P_i to a path P_i' with the same endpoints as Q_1, by adding to it at most two outer edges of $\hat{S}_i$. But P_i' is then shorter than both Q_1 and Q_2, in contradiction to the fact that C_i is ℓ_i-geodetic. Thus there is no shortcut to C and therefore it is ℓ-geodetic. $\square$

4.3 Proof of the generating theorem

Before we are able to prove Theorem 1.2, we need one last lemma.

Lemma 4.11. *Let C be a circle in $|G|$ and let $i \in \mathbb{N}$ be minimal such that C meets S_i. Then, there exists a finite family $\mathcal{F}$, each element of which is an ℓ-geodetic circle in $|G|$ of length at most $5\varepsilon_i$, and such that $\sum \mathcal{F}$ coincides with C in $\tilde{S}_i$, that is, $(\sum \mathcal{F}) \cap \tilde{S}_i = C \cap \tilde{S}_i$.*

Proof. For every $j \geqslant i$, choose, among all families $\mathcal{H}$ of ℓ_j-geodetic cycles in $\hat{S}_j$, the ones that are minimal with the following properties, and let V_j be their set:

- no cycle in $\mathcal{H}$ is longer than $5\varepsilon_i$ in $\hat{S}_j$, and
- $\sum \mathcal{H}$ coincides with C in $\tilde{S}_i$.

Note that every cycle in such a family $\mathcal{H}$ meets S_i as otherwise $\mathcal{H}$ would not be minimal with the above properties. By Lemma 4.3, the sets V_j are not empty. As no family in V_j contains a cycle twice, and $\hat{S}_j$ has only finitely many cycles, every V_j is finite. Our aim is to find a family $\mathcal{C}_i \in V_i$ so that each cycle in $\mathcal{C}_i$ can be extended to an ℓ-geodetic circle, giving us the desired family $\mathcal{F}$.

Given $j \geqslant i$ and $\mathcal{C} \in V_{j+1}$, restricting every cycle in $\mathcal{C}$ to $\hat{S}_j$ (note that by the minimality of $\mathcal{C}$, every cycle in $\mathcal{C}$ meets S_i and hence also its supset S_j) yields, by Lemma 4.9, a family $\mathcal{C}^-$ of ℓ_j-geodetic cycles. Moreover, $\mathcal{C}^-$ lies in V_j: by Lemma 4.1, no element of $\mathcal{C}^-$ is longer than $5\varepsilon_i$, and the sum of $\mathcal{C}^-$ coincides with C in $\hat{S}_i$, as the performed restrictions do not affect the edges in $\tilde{S}_i$. In addition, $\mathcal{C}^-$ is minimal with respect to the above properties as $\mathcal{C}$ is.

Now construct an auxiliary graph with vertex set $\bigcup_{j \geqslant i} V_j$, where for every $j > i$, every element $\mathcal{C}$ of V_j is incident with $\mathcal{C}^-$. Applying Lemma 2.7 to this graph, we obtain an infinite sequence $\mathcal{C}_i, \mathcal{C}_{i+1}, \ldots$ such that for every $j \geqslant i$, $\mathcal{C}_j \in V_j$ and $\mathcal{C}_j = \mathcal{C}_{j+1}^-$. Therefore, for every cycle $C \in \mathcal{C}_j$ there is a unique cycle in $\mathcal{C}_{j+1}$ whose restriction to $\hat{S}_j$ is C. Hence for every $D \in \mathcal{C}_i$ there is a chain $(D =)D_i, D_{i+1}, \ldots$ of cycles such that $D_j \in \mathcal{C}_j$ for every $j \geqslant i$. By Lemma 4.10, the limit X_D of this chain is an ℓ-geodetic circle, and X_D is not longer than $5\varepsilon_i$, because in that case some D_j would also be longer than $5\varepsilon_i$. Thus, the family $\mathcal{F}$ resulting from $\mathcal{C}_i$ after replacing each $D \in \mathcal{C}_i$ with X_D has the desired properties. $\square$

Proof of Theorem 1.2. If $(\mathcal{F}_i)_{i \in I}$ is a family of families, then let the family $\bigcup_{i \in I} \mathcal{F}_i$ be the disjoint union of the families $\mathcal{F}_i$. Let C be an element of $\mathcal{C}(G)$. For $i = 0, 1, \ldots$, we define finite families Γ_i of ℓ-geodetic circles that satisfy the following condition:

$$C_i := C + \sum \bigcup_{j \leqslant i} \Gamma_j \text{ does not contain edges of } \tilde{S}_i, \tag{2}$$

where $+$ denotes the symmetric difference.

By Lemma 2.2, there is a family $\mathcal{C}$ of edge-disjoint circles whose sum equals C. Applying Lemma 4.11 to every circle in $\mathcal{C}$ that meets S_0 (there are only finitely many), yields a finite family Γ_0 of ℓ-geodetic circles that satisfies condition (2).

Now recursively, for $i = 0, 1, \ldots$, suppose that $\Gamma_0, \ldots, \Gamma_i$ are already defined finite families of ℓ-geodetic circles satisfying condition (2), and write C_i as a sum of a family $\mathcal{C}$ of edge-disjoint circles, supplied by Lemma 2.2. Note that only finitely many members of $\mathcal{C}$ meet S_{i+1}, and they all avoid S_i as C_i does. Therefore, for every member D of $\mathcal{C}$ that meets S_{i+1}, Lemma 4.11 yields a finite family $\mathcal{F}_D$ of ℓ-geodetic circles of length at most $5\varepsilon_{i+1}$ such that $(\sum \mathcal{F}_D) \cap \tilde{S}_{i+1} = D \cap \tilde{S}_{i+1}$. Let

$$\Gamma_{i+1} := \bigcup_{\substack{D \in \mathcal{C} \\ D \cap S_{i+1} \neq \emptyset}} \mathcal{F}_D.$$

By the definition of C_i and Γ_{i+1}, we have

$$C_{i+1} = C + \sum \bigcup_{j \leqslant i+1} \Gamma_j = C + \sum \bigcup_{j \leqslant i} \Gamma_j + \sum \Gamma_{i+1} = C_i + \sum \Gamma_{i+1}.$$

By the definition of Γ_{i+1}, condition (2) is satisfied by C_{i+1} as it is satisfied by C_i. Finally, let

$$\Gamma := \bigcup_{i<\omega} \Gamma_i.$$

Our aim is to prove that $\sum \Gamma = C$, so let us first show that Γ is thin.

We claim that for every edge $e \in E(G)$, there is an $i \in \mathbb{N}$, such that for every $j \geqslant i$ no circle in Γ_j contains e. Indeed, there is an $i \in \mathbb{N}$, such that ε_j is smaller than $\frac{1}{5}\ell(e)$ for every $j \geqslant i$. Thus, by the definition of the families Γ_j, for every $j \geqslant i$, every circle in Γ_j is shorter than $\ell(e)$, and therefore too short to contain e. This proves our claim, which, as every Γ_i is finite, implies that Γ is thin.

Thus, $\sum \Gamma$ is well defined; it remains to show that it equals C. To this end, let e be any edge of G. By (2) and the claim above, there is an i, such that e is contained neither in C_i nor in a circle in $\bigcup_{j>i} \Gamma_i$. Thus, we have

$$e \notin C_i + \sum \bigcup_{j>i} \Gamma_j = C + \sum \Gamma.$$

As this holds for every edge e, we deduce that $C + \sum \Gamma = \emptyset$, so C is the sum of the family Γ of ℓ-geodetic circles. $\square$

5 Further problems

It is known that the finite circles (i.e. those containing only finitely many edges) of a locally finite graph G generate $\mathcal{C}(G)$ (see [7, Corollary 8.5.9]). In the light of this result and Theorem 1.2, it is natural to pose the following question:

Problem 1. *Let G be a locally finite graph, and consider a metric representation $(|G|, \ell)$ of G. Do the finite ℓ-geodetic circles generate $\mathcal{C}(G)$?*

The answer to Problem 1 is negative: Figure 5.5 shows a graph with a metric representation where no geodetic circle is finite.

In Section 2 we did not define $d_\ell(x, y)$ as the length of a shortest x–y arc, because we could not guarantee that such an arc exists. But does it? The following result asserts that it does.

Proposition 5.1. *Let a metric representation $(|G|, \ell)$ of a locally finite graph G be given and write $d = d_\ell$. For any two distinct points $x, y \in \hat{V}$, there exists an x–y arc in $|G|$ of length $d(x, y)$.*

Proof. Let $\mathcal{P} = P_0, P_1, \ldots$ be a sequence of x–y arcs in $|G|$ such that $(\ell(P_j))_{j\in\mathbb{N}}$ converges to $d(x, y)$. Choose a $j \in \mathbb{N}$ such that every arc in $\mathcal{P}$ meets S_j, where S_j, $\tilde{S}_j$, and $\hat{S}_j$ are defined as before. Such a j always exists; if for example $x, y \in \Omega$, then pick j so that S_j separates a ray in x from a ray in y.

As $\hat{S}_j$ is finite, there is a path X_j in $\hat{S}_j$ and a subsequence $\mathcal{P}_j$ of $\mathcal{P}$ such that X_j is the restriction of any arc in $\mathcal{P}_j$ to $\hat{S}_j$. Similarly, for every $i > j$, we can recursively find a path X_i in $\hat{S}_i$ and a subsequence $\mathcal{P}_i$ of $\mathcal{P}_{i-1}$ such that X_i is the restriction of any arc in $\mathcal{P}_i$ to $\hat{S}_i$.

By construction, $X_j, X_{j+1}, \ldots$ is a chain of paths. The limit X of this chain contains x and y as it is closed, and $\ell(X) \leqslant d(x,y)$; for if $\ell(X) > d(x,y)$, then there is an i such that $\ell(X_i \cap \tilde{S}_i) > d(x,y)$, and as $\ell(X_k \cap \tilde{S}_k) \geqslant \ell(X_i \cap \tilde{S}_i)$ for $k > i$, this contradicts the fact that $(\ell(P_j))_{j\in\mathbb{N}}$ converges to $d(x,y)$. By Lemma 4.7, X is the image of a topological path and thus, by Lemma 2.5, contains an x–y arc P. Since P is at most as long as X, it has length $d(x,y)$ (thus as $\ell(X) \leqslant d(x,y)$, we have $P = X$.) $\square$

Our next problem raises the question of whether it is possible, given $x, y \in V(G)$, to approximate $d(x,y)$ by finite x–y paths:

Problem 2. *Let G be a locally finite graph, consider a metric representation $(|G|, \ell)$ of G and write $d = d_\ell$. Given $x, y \in V(G)$ and $\epsilon \in \mathbb{R}_+$, is it always possible to find a finite x–y path P such that $\ell(P) - d(x,y) < \epsilon$?*

Surprisingly, the answer to this problem is also negative. The graph of Figure 5.5 with the indicated metric representation is again a counterexample.

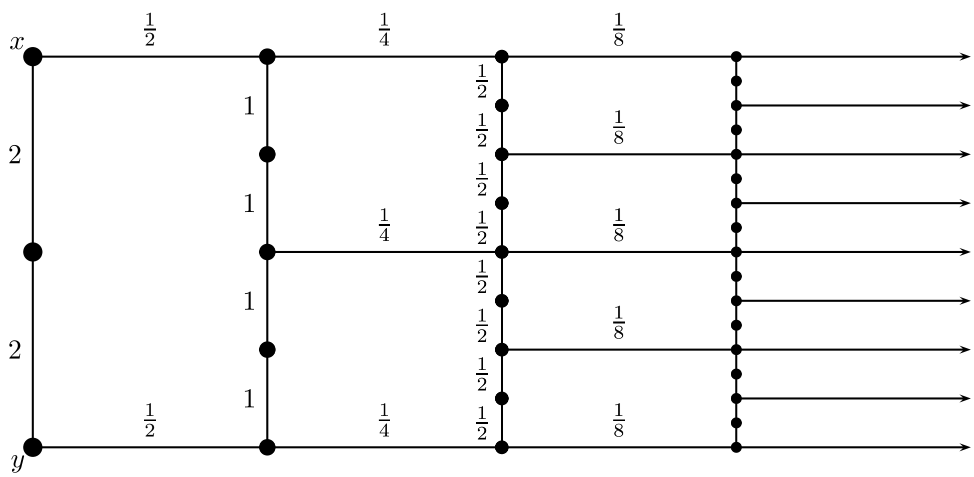

Fig. 5.5: A 1-ended graph G with a metric representation. Every ℓ-geodetic circle is easily seen to contain infinitely many edges. Moreover, every (graph-theoretical) x–y path has length at least 4, although $d(x,y) = 2$.

As noted in Section 4, every ℓ-geodetic circle has finite length. But what about other circles? Is it possible to choose a metric representation such that there are circles of infinite length? Yes it is, Figure 5.5 shows such a metric representation. It is even possible to have every infinite circle have infinite length: Let G be the infinite ladder, let the edges of the upper ray have lengths $\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots$, let the edges of the lower ray have lengths $\frac{1}{2}, \frac{1}{3}, \frac{1}{4}, \ldots$, and let the rungs have lengths $\frac{1}{2}, \frac{1}{4}, \frac{1}{8}, \ldots$. This clearly yields a metric representation, and as any infinite circle contains a tail of the lower ray, it has infinite length. This means that in this example all ℓ-geodetic circles are finite, contrary to the metric representation in Figure 5.5, where every ℓ-geodetic circle is infinite.

Theorems 3.1 and 1.2 could be applied in order to prove that the cycle space of a graph is generated by certain subsets of its, by choosing an appropriate length function, as indicated by our next problem. Call a cycle in a finite graph *peripheral*, if it is induced and non-separating.

Problem 3. *If G is a 3-connected finite graph, is there an assignment of lengths ℓ to the edges of G, such that every ℓ-geodetic cycle is peripheral?*

We were not able to give an answer to this problem. A positive answer would imply, by Theorem 3.1, a classic theorem of Tutte [18], asserting that the peripheral cycles of a 3-connected finite graph generate its cycle space. Problem 3 can also be posed for infinite graphs, using the infinite counterparts of the concepts involved[2].

Acknowledgement

The counterexamples in Figures 3.1 and 5.5 are due to Henning Bruhn. The authors are indebted to him for these counterexamples and other useful ideas.

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

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

[8] R. Diestel. End spaces and spanning trees. *J. Combin. Theory (Series B)*, 96:846–854, 2006.

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

[2]Tutte's theorem has already been extended to locally finite graphs by Bruhn [2]

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

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

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

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

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

[15] A. Georgakopoulos. Topological circles and Euler tours in locally finite graphs. *Electronic J. Comb.*, 16:#R40, 2009.

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

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

[18] W.T. Tutte. How to draw a graph. *Proc. London Math. Soc.*, 13:743–768, 1963.

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