


Topological paths and cycles in infinite graphs

Dissertation

zur Erlangung des Doktorgrades
der Fakultät für Mathematik, Informatik
und Naturwissenschaften
der Universität Hamburg



vorgelegt
im Department Mathematik
von
Angelos Georgakopoulos
aus Athen

Hamburg 2007

Als Dissertation angenommen vom Department
Mathematik der Universität Hamburg

auf Grund der Gutachten von Prof. Reinhard Diestel, PhD
und Prof. Dr. Thomas Andreae

Hamburg, den 20. Dezember 2006

Prof. Dr. Hans Joachim Oberle
Leiter des Departments Mathematik

Contents

1	Overview	1
2	Definitions and basic facts	5
2.1	Definitions	5
2.2	Basic facts	6
2.3	Homeomorphisms between the end-space of a graph and a subgraph	8
3	Warming-up: The cube of a locally finite graph is hamiltonian	11
4	Connectedness vs. path-connectedness in G	15
4.1	Introduction	15
4.2	Connectedness does not imply path-connectedness in $ G $. . .	16
4.3	Connectedness implies path-connectedness almost always . . .	21
4.4	Connectedness implies path-connectedness for closed subspaces	27
5	Linear algebra with infinite sums	31
5.1	Introduction	31
5.2	Bases	32
5.3	Closedness	37
6	Geodesic circles	43
6.1	Introduction	43
6.2	Generating $\mathcal{C}(\mathbf{G})$ by geodesic cycles	44
6.2.1	Finite graphs	44
6.2.2	Failure in infinite graphs	45
6.3	Generating $\mathcal{C}(\mathbf{G})$ by geodesic circles	46
6.3.1	Restricting paths and circles	46
6.3.2	Limits of paths and cycles	49
6.3.3	Proof of the generating theorem	52

6.4	Further problems	54
7	Hamilton circles in squares of locally finite graphs	57
7.1	Introduction	57
7.2	Finite graphs	58
7.2.1	Definitions	58
7.2.2	Proof of Fleischner's theorem	58
7.2.3	Total graphs	61
7.3	Infinite graphs	64
7.3.1	Definitions	64
7.3.2	Outline of the proof of Theorem 7.2	65
7.3.3	Some preliminaries	70
7.4	Proof of Theorem 7.2	73
7.4.1	A stronger assertion	73
7.4.2	Constructing the scaffolding	74
7.4.3	Making the graph eulerian	80
7.4.4	Splitting into finite multigraphs	82
7.4.5	Cleaning up the articulation points	93
7.4.6	The hamiltonisation	95
7.5	Final remarks	96
	Bibliography	99

Acknowledgement

I am very grateful to my advisor Reinhard Diestel for his invaluable help in mathematics and everyday matters. I would also like to thank the University of Hamburg for supporting me with scholarships.

Chapter 1

Overview

This thesis is about infinite graphs. Its main result is the extension to infinite, locally finite graphs of a well known theorem of Fleischner about the square of a finite graph. The n -th power G^n of a graph G is the graph on $V(G)$ in which two vertices are adjacent if and only if they have distance at most n in G . Fleischner's theorem states that

Theorem 1.1 (Fleischner [17, 18]). *If G is a finite 2-connected graph, then G^2 is Hamiltonian.*

Settling a conjecture of Diestel [11, 12] we will, in Chapter 7, fully extend this fact to locally finite graphs:

Theorem 1.2 ([21]). *If G is a locally finite 2-connected graph, then G^2 is Hamiltonian.*

An extension of Theorem 1.1 to infinite graphs had already been proved by Thomassen [29], for the special case of those locally finite graphs in which the removal of any finite set of vertices leaves precisely one infinite component behind. Fleischner's original proof of Theorem 1.1 was long and complicated. Using Thomassen's method, Říha (see [32] or [12]) produced a shorter proof, which was still quite long. Interestingly, the study of infinite graphs led once more to a shorter proof of the same theorem: in Chapter 7 we will see a new proof of Theorem 1.1 shorter than that of Říha, which resulted from an idea used in the proof of Theorem 1.2.

We stated Theorem 1.2 without mentioning what it means for an infinite graph to be Hamiltonian. In fact, this is hard to define; it is easy to visualise an infinite path: an infinite sequence of vertices such that each of them is connected to the next by an edge. But infinite cycles — let alone infinite Hamilton cycles — cannot be defined that way, because an infinite sequence cannot come back to its starting point. There is however an elegant and

successful definition of infinite cycles. It is a concept called *circle* recently invented by Diestel and Kühn (see [11] or [12]), and it has the properties one would expect from an infinite cycle: every vertex in it has precisely two neighbours, and it is “round” in the sense that it comes back to its starting point. A circle is a homeomorphic image of S^1 , the unit circle in \mathbb{R}^2 , in a topological space $|G|$ that consists of a locally finite graph G seen as a 1-complex and some extra points called the *ends of G* , which can be thought of as points at infinity (see Chapter 2). A *Hamilton circle* is, then, a circle containing all vertices of a graph, and a locally finite graph is *Hamiltonian* if it has a Hamilton circle.

Similarly to Theorem 1.2, we also extend to locally finite graphs the fact that the third power of a connected finite graph has a Hamilton cycle:

Theorem 1.3 ([21]). *If G is a connected locally finite graph, then G^3 has a Hamilton circle.*

Another attempt to use the concept of Hamilton circle to generalise a finite theorem has been made by Bruhn and Yu [8], who partly generalised Tutte’s Theorem [30] that a finite 4-connected planar graph has a Hamilton cycle.

The notion of circle has been coupled by another new notion, that of the (*topological*) *cycle space* of an infinite graph, to produce a very strong tool in infinite graph theory. The cycle space of a finite graph G is a vector space over \mathbb{Z}_2 , whose elements are those subsets of $E(G)$ that can be obtained as sums, with symmetric difference as addition, of edge sets of cycles of G . In order to make this concept suitable for infinite graphs, it has to be extended in two ways: on the one hand, we have to allow as elements edge sets of circles in addition to edge sets of finite cycles, and on the other, we have to allow certain sums (with symmetric difference as addition) of infinitely many summands. The infinite sums we allow are the ones for which the family of summands is *thin*, that is, no edge lies in infinitely many of its elements. These are precisely the families of summands for which we can decide for each edge if it lies in an odd or even number of summands. The sum of a thin family of edge sets is, then, the set of those edges that lie in an odd number of elements of the family (see Chapter 2 for precise definitions and [11] for more).

There has been a number of results recently by Bruhn Diestel Kühn and Stein that exploit the concepts of circle and topological cycle space in order to extend well known results about finite graphs to infinite ones, as Theorem 1.2 does. Among these results are: Euler’s theorem that a graph is Eulerian if and only if every vertex has even degree [7, 14], MacLane’s planarity criterion [6], Tutte’s theorem that the peripheral cycles of a 3-connected finite graph

generate its cycle space and Tutte's planarity criterion [2], Whitney's theorem that a finite graph has a dual if and only if it is planar [3], Gallai's theorem that every finite graph has a vertex partition into two parts each inducing an element of its cycle space [4], the Tutte/Nash-Williams packing theorem [11], and several basic facts about the cycle space of a finite graph [11].

In Chapter 6 we will give a result of this kind. A finite cycle C in a graph G is called *geodesic* if, for any two vertices $x, y \in C$, the length of at least one of the two x - y paths on C equals the distance of x and y in G . It is easy to prove that (see Chapter 6):

Theorem 1.4 ([23]). *The cycle space of a finite graph is generated by the circuits of its geodesic cycles.*

In Chapter 6 we show that seen in the right setting, Theorem 1.4 generalises to locally finite graphs.

The theory of circles and the related cycle space does not only help to generalise well known facts about finite graphs, but also raises questions that reach beyond finite graph theory and are interesting in their own right. In Chapter 5 we will see such a case: stimulated by problems concerning the cycle space of an infinite graph, where as already mentioned sums of infinitely many edge sets are allowed, we will study the ramifications of allowing sums of infinite, thin families of summands in vector spaces and modules in general. The questions we will pose are whether every generating set in such a setting contains a *basis*, i.e. a minimal subset generating the same space as the original set, and whether the space generated by some set is closed under taking infinite sums. The answers we will give are applicable in infinite graph theory but also interesting from the algebraic point of view.

A further problem that is interesting for infinite graphs but not for finite ones was posed by Diestel and Kühn, who conjectured [16, 11] that if G is a locally finite graph then every connected subspace of $|G|$ is path-connected. As discussed in Chapter 4, this problem is important for many applications, because reducing path-connectedness to connectedness can facilitate the proof that a certain subspace of $|G|$ is a circle, as circles are by definition path-connected. In Chapter 4 we settle this conjecture negatively, by constructing a counterexample:

Theorem 1.5 ([20]). *There exists a locally finite graph G such that $|G|$ has a connected subspace which is not path-connected.*

This thesis is structured as follows. Chapter 2 contains definitions and some basic facts used more or less in the whole thesis. In Chapter 3 we prove Theorem 1.3, and the reader is encouraged to read this proof before

going on to the rest of the thesis, especially Chapter 7, as it exemplifies the most important tools used later. An exception is Chapter 5, where we study the consequences of allowing infinite sums in abstract modules, which can be read independently. Chapter 4 discusses and proves Theorem 1.5, and Chapter 6 Theorem 1.4. Finally, in Chapter 7 we prove Theorem 1.2.

Chapter 2

Definitions and basic facts

2.1 Definitions

We are using the terminology of [12] for graph theoretical concepts, that of [1] for topological concepts and that of [9] for logical ones.

Let $G = (V, E)$ be a *locally finite* graph — i.e. every vertex has a finite degree — fixed throughout this section. A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*. Two rays R, L in G are *equivalent* if no finite set of vertices separates them; we denote this fact by $R \approx_G L$, or simply by $R \approx L$ if G is fixed. The corresponding equivalence classes of rays are the *ends* of G . We denote the set of these ends by $\Omega = \Omega(G)$.

Let G bear the topology of a 1-complex¹. 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 subray 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

$$C(S, \omega) \cup \Omega(S, \omega) \cup E'(S, \omega) \tag{2.1}$$

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 . For any given such ω and S , pick one of these sets and denote it by $O(S, \omega)$. Let $|G|$ denote the topological space of $G \cup \Omega$ endowed with the topology generated by the open sets of the form (2.1) together with those of the 1-complex G .

¹Every 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 .

It can be proved (see [13]) that in fact $|G|$ is the Freudenthal compactification [19] of the 1-complex G .

An inner point of an edge of the 1-complex G will be called an *edge point*.

For any vertex $v \in V$ let $N^i(v)$ denote the set of vertices of G whose distance from v is at most i (including v), and let $G[v]^i$ be the subgraph of G induced by $N^i(v)$.

A continuous map σ from the real unit interval $[0, 1]$ to a topological space X is a (topological) *path* in X ; the images under σ of 0 and 1 are its *endpoints*. A homeomorphic image (in the subspace topology) of $[0, 1]$ in a topological space X will be called an *arc* in X . Any set $\{x\}$ with $x \in |G|$ is also called an arc in $|G|$.

We now define one of the central concepts of this thesis: a (*topological cycle* or) *circle* in $|G|$ is a homeomorphic image of S^1 , the unit circle in \mathbb{R}^2 , in $|G|$. A *Hamilton circle of G* is a circle that contains every vertex of G (and hence, also every end, as it is closed).

A subset D of E is a *circuit* if there is a circle C in $|G|$ such that $D = \{e \in E | e \subseteq C\}$. Call a family $\mathcal{F} = (D_i)_{i \in I}$ of subsets of E *thin*, if no edge lies in D_i for infinitely many indices i . Let the *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 .

We can now define a further central concept of this thesis: let the (*topological*) *cycle space* $\mathcal{C}(G)$ of G be the set of all sums of (thin families of) circuits.

A *normal* spanning tree of G is a spanning tree T of G with a root r such that any two adjacent vertices in G are comparable in the tree-order of T .

2.2 Basic facts

The following two lemmas are perhaps the most fundamental facts about the cycle space of an infinite graph. Both can be found in [12, Theorem 8.5.8]. Let G be an arbitrary connected locally finite multigraph fixed throughout this section (these results were proved for simple graphs, but their generalisation to multigraphs is trivial).

Lemma 2.1. *Every element of $\mathcal{C}(G)$ is a disjoint union of circuits.*

Lemma 2.2. *Let $F \subseteq E(G)$. Then $F \in \mathcal{C}(G)$ if and only if F meets every finite cut in an even number of edges.*

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

Lemma 2.3 ([12, Proposition 8.5.1]). *If G is locally finite and connected, then $|G|$ is a compact Hausdorff space.*

We will say that a topological path *traverses* an edge xy if it maps an interval of $[0, 1]$ onto xy .

Lemma 2.4 ([20]). *Any topological path that connects some point p of a basic open neighbourhood U of an end to a point q outside U must traverse some edge xy with $x \in U, y \notin U$, unless p or q lies in such an edge xy .*

Proof. Let R be the image of such a path, and suppose it avoids all edges between U and $V(G) \setminus U$ (It is easy to see that, without loss of generality, R either traverses any given edge xy or does not meet (x, y) at all, unless p or q lie in xy). Then both $U \cap R$ and $(V(G) \setminus U) \cap R$ are open in the subspace topology of R , which shows that R is disconnected. But this cannot be true since R is a continuous image of $[0, 1]$. \square

The following basic topological lemma can be found in [25, 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 .*

The union of a ray R with infinitely many disjoint finite paths having precisely their first vertex on R is a *comb*; the last vertices of those paths are the *teeth* of this comb, and R is its *spine*. The following very basic lemma can be found in [12, 8.2.2].

Lemma 2.6. *If U is an infinite set of vertices in G , then G contains a comb with all teeth in U .*

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

Lemma 2.7 (König's Infinity Lemma [27]). *Let V_0, V_1, \dots 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 $n \geq 1$ has a neighbour in V_{n-1} . Then G contains a ray $v_0v_1 \dots$ with $v_n \in V_n$ for all n .*

König's Infinity Lemma is closely related with the compactness theorem for propositional logic, which we will also use:

Theorem 2.1 (Compactness Theorem [9]). *Let K be an infinite set of propositional formulas, every finite subset of which is satisfiable. Then K is satisfiable.*

2.3 Homeomorphisms between the end-space of a graph and a subgraph

The results of this section will be used in Chapter 3 and Chapter 7.

If H is a spanning subgraph of some graph G , then there is usually no need to distinguish between vertices of H and vertices of G . For ends however, the matters are more complicated. In what follows, we develop some tools that will in some cases help us work with the ends of H as if they were the ends of G .

For any two multigraphs $H \subseteq G$, define the mapping π_{HG} by

$$\begin{aligned} \pi_{HG} : \Omega(H) &\rightarrow \Omega(G) \\ \omega &\mapsto \omega' \supseteq \omega \end{aligned}$$

Lemma 2.8. *Let H, G be locally finite connected multigraphs such that $H \subseteq G$, $V(H) = V(G)$, and for any two rays R, S in H , if $R \approx_G S$ then $R \approx_H S$. Then π_{HG} is a homeomorphism between $\Omega(H)$ and $\Omega(G)$.*

Proof. Clearly, π_{HG} is injective. Let us show that it is surjective. For any $\omega \in \Omega(G)$, pick a ray $R \in \omega$. Since H is connected, we can apply Lemma 2.6 to obtain a comb in H with teeth in $V(R)$. The spine of this comb is a ray in H , that is equivalent to R in G . Thus its end is mapped to ω by π_{HG} .

Since $H \subseteq G$, it follows easily that π_{HG} is continuous. Moreover, $\Omega(H)$ is compact, because it is closed in $|H|$ and $|H|$ is compact by Lemma 2.3. It is an elementary topological fact ([1, Theorem 3.7]) that a continuous bijection from a compact space to a Hausdorff space is a homeomorphism, which implies that π_{HG} is indeed a homeomorphism between $\Omega(H)$ and $\Omega(G)$. \square

Lemma 2.9. *Let H, G be locally finite connected multigraphs such that $H \subseteq G$, $V(H) = V(G)$, and for any two rays R, S in H , if $R \approx_G S$ then $R \approx_H S$. Let $(v_i)_{i \in \mathbb{N}}$ be a sequence of vertices of $V(G)$. Then v_i converges to $\omega \in \Omega(H)$ in $|H|$ if and only if v_i converges to $\pi_{HG}(\omega)$ in $|G|$.*

Proof. Define a mapping $\hat{\pi}_{HG} : V(H) \cup \Omega(H) \rightarrow V(G) \cup \Omega(G)$ that maps every end $\omega \in \Omega(H)$ to $\pi_{HG}(\omega)$, and every vertex in $V(H)$ to itself. Easily by Lemma 2.8, $\hat{\pi}_{HG}$ is bijective and continuous. Moreover, $V(H) \cup \Omega(H)$ is closed, thus compact, so like in the proof of Lemma 2.8, $\hat{\pi}_{HG}$ is a homeomorphism between $V(H) \cup \Omega(H)$ and $V(G) \cup \Omega(G)$, from which the assertion easily follows. \square

For any two connected multigraphs G, H such that $V(G) = V(H)$, we will write $|H| \cong |G|$ if there is a homeomorphism $\pi : \Omega(H) \rightarrow \Omega(G)$, such

that for any sequence $(v_i)_{i \in \mathbb{N}}$ of vertices of $V(G)$, v_i converges to $\omega \in \Omega(H)$ in $|H|$ if and only if v_i converges to $\pi(\omega)$ in $|G|$.

If $H \subseteq G$, and $e = uv \in E(G) - E(H)$, then a *detour for e* (in H) is a path in H with endvertices u, v .

Lemma 2.10. *Let $H \subseteq G$ be locally finite multigraphs such that $V(H) = V(G)$ and G is connected. Suppose that for each edge $e \in E(G) - E(H)$, a detour $dt(e)$ for e has been specified. If the set $\{dt(e) | e \in E(G) - E(H)\}$ is thin, i.e. no edge appears in infinitely many of its elements, then $|H| \cong |G|$.*

Proof. Clearly, H is connected. Pick any two rays R, S in H , such that $R \approx_G S$. By Lemmas 2.8 and 2.9, it suffices to show that $R \approx_H S$.

Since $R \approx_G S$, there is an infinite set \mathcal{P} of disjoint R - S -paths in G . For each $P \in \mathcal{P}$, replace all edges e of P not in $E(H)$ with $dt(e)$, to obtain a connected subgraph P' of H containing the endvertices of P . Let $dt(P)$ be an R - S -path in P' . The set of all these paths $\{dt(P) | P \in \mathcal{P}\}$ is clearly thin, proving that $R \approx_H S$. \square

Chapter 3

Warming-up: The cube of a locally finite graph is hamiltonian

Let us start the main part of this thesis with an easy but non-trivial result whose proof makes use of the most fundamental results and methods in the study of $|G|$ and $\mathcal{C}(G)$, and is thus appropriate for preparing the reader for the more difficult results of Chapter 4 and especially of Chapter 7:

Theorem 3.1 ([21]). *If G is a connected locally finite graph, then G^3 has a Hamilton circle.*

For finite graphs this is well known. Extensions to infinite graphs had already been made by Sekanina [28], who showed that the third power of a connected, locally finite, 1-ended graph has a spanning ray, and by Heinrich [26], who specified a class of non-locally-finite graphs, whose third power has a spanning ray. With Theorem 3.1, which we will now prove, we generalise to locally finite graphs with any number of ends.

Proof of Theorem 3.1. We will say that an edge $e = uv$ of some graph G crosses a subgraph H of G , if $u \in V(H)$ and $v \notin V(H)$. An x -branch of a tree T with root v , for some vertex $x \in V(T)$, is a component of $T - x$ that does not contain v ; a subgraph of T is a branch, if it is an x -branch for some $x \in V(T)$.

Let T be a normal spanning tree of G , with root v (every countable connected graph has a normal spanning tree, see [12, Theorem 8.2.4]), and let $T_i = T[v]^i$.

We will prove the assertion using Theorem 2.1. To this end, define for each edge $e \in E(T^3)$ a logical variable $v(e)$, the truth-values of which encode

the presence or not of e , and let \mathcal{V} be the set of these variables. For every vertex $x \in V(G)$, write a propositional formula with variables in \mathcal{V} , expressing the fact that exactly two x -edges are present, and let \mathcal{P}_1 be the set of these formulas. For every branch B of T , write a propositional formula with variables in \mathcal{V} expressing the fact that at most two edges that cross B are present, and let \mathcal{P}_2 be the set of these formulas. For every finite cut F of T^3 , write a propositional formula with variables in \mathcal{V} , expressing the fact that an even, positive number of edges in F are present, and let \mathcal{P}_3 be the set of these formulas. Let $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2 \cup \mathcal{P}_3$.

For every finite $\mathcal{P}' \subseteq \mathcal{P}$, there is an assignment of truth-values to the elements of \mathcal{V} , satisfying all elements of \mathcal{P}' : if i is large enough, then by the following lemma, T_i^3 has a Hamilton cycle, which encodes such an assignment:

Lemma 3.1. *If T is a finite tree with root v and $|T| \geq 3$, then T^3 has a Hamilton cycle H , that contains a v -edge $e(H) \in E(T)$, and for every branch B of T , H contains precisely two edges that cross B .*

Proof (sketch). We will use induction on the height h of T . The assertion is clearly true for $h = 1$. If $h > 1$, then apply the induction hypothesis on each non-trivial v -branch, delete $e(H_v)$ for each resulting Hamilton cycle H_v , and use some edges of T^3 as shown in Figure 3.1, to construct the desired Hamilton cycle H of T^3 . It is easy to see that no branch of T is crossed by more than two edges of H , if this is true for the Hamilton cycles H_v of the v -branches. \square




Figure 3.1: Using the induction hypothesis to pick a Hamilton cycle of T . The wavy curves represent Hamilton cycles of the v -branches supplied by the induction hypothesis, and for each such Hamilton cycle H , $e(H)$ is represented by a dashed line. The thick cycle represents H .

So by Theorem 2.1, there is an assignment of truth-values to the elements of \mathcal{V} , satisfying all elements of \mathcal{P} . Let F be the set of edges that are

present according to this assignment. We will prove that F is the circuit of a Hamilton circle of T^3 .

By Lemma 2.2, $F \in \mathcal{C}(T)$, thus by Lemma 2.1, F is a disjoint union of circuits. Let $C \subseteq F$ be a circuit, and suppose, for contradiction, that there is a vertex $u \in T$ not incident with C . Choose an $i \in \mathbb{N}$ so that T_i meets both u and C . If $V(C) \subseteq V(T_i)$, then $V(C)$ defines a finite cut, which is not met by F , because otherwise a formula in \mathcal{P}_1 is contradicted; this, however, contradicts a formula in \mathcal{P}_3 . If $V(C) \not\subseteq V(T_i)$, let \mathcal{B} be the (non-empty) set of branches B in $T - T_i$ such that $B \cap C \neq \emptyset$, and let $X = V(C) \cup \bigcup_{B \in \mathcal{B}} V(B)$. Since $u \notin X$, $E(X, X' := V(T) - X)$ is a non-empty cut D , which is clearly finite. Now for every $B \in \mathcal{B}$, there is a formula in \mathcal{P}_2 asserting that there are at most two edges crossing B , and since (by Lemma 2.4 and Lemma 2.2) C already contains two such edges, F contains no X' - B -edge. Moreover, F contains no X' - C -edge, because of the formulas in \mathcal{P}_1 , thus $D \cap F = \emptyset$, contradicting a formula in \mathcal{P}_3 .

Thus F is the circuit of a Hamilton circle H of T^3 . Applying Lemma 2.10 on T, T^3 , using a path of length at most 3 as a detour for each edge in $E(T^3) - E(T)$, we obtain $|T^3| \cong |T|$, and similarly $|G^3| \cong |G|$. Easily by Lemma 2.9, $|T| \cong |G|$, thus H is also a Hamilton circle of G^3 . \square

Chapter 4

Connectedness vs. path-connectedness in $|G|$

4.1 Introduction

In this chapter we give an answer to the following question, where G is a locally finite graph fixed throughout this section:

Problem 4.1 ([11]). *Is every connected subspace of $|G|$ path-connected?*

Apart from being interesting as a basic topological question in its own right, this problem is also important from the graph-theoretical point of view. Indeed, in order to prove that a certain subspace of $|G|$ is a circle or a topological spanning tree, one has to show that it is path-connected, but it is often much easier to show that it is (topologically) connected. See for example Theorem 8.5.8 in [12], which summarizes the basic properties of the cycle space of a locally finite graph. The reduction of path-connectedness to connectedness simplifies its proof considerably in comparison to the original proof in [15, Theorem 5.2]. Lemma 8.5.13 in [12] is an example of how reducing path-connectedness to connectedness can facilitate proving the existence of a topological spanning tree, which can otherwise be a tedious task as witnessed by the proof of Theorem 5.2 in [16]. Further examples include Exercises 65 and 70 of [12], which describe some fundamental properties of circles and topological spanning trees, and Lemma 6.6 in Chapter 6: their proofs become easy when the path-connectedness required is replaced with connectedness, while without this tool they would be arduous and long.

Diestel and Kühn [16] have shown that:

Theorem 4.1. *Every closed connected subspace of $|G|$ is path-connected.*

In Section 4.4 we give an alternative proof to Theorem 4.1. It was conjectured [16, 11] that the answer to Problem 4.1 should be positive also in general. However, we shall construct a counterexample (Section 4.2):

Theorem 4.2 ([20]). *There exists a locally finite graph G such that $|G|$ has a connected subspace which is not path-connected.*

The counterexample has a complicated structure, but as we shall see in Sections 4.3 and 4.4 every counterexample to Problem 4.1 has to be that complicated. In particular, some restrictions on the structure of such a counterexample are posed by the following result (proved in Section 4.3):

Theorem 4.3 ([20]). *Given any locally finite connected graph G , a connected subspace X of $|G|$ is path-connected unless it satisfies the following assertions:*

- *X has uncountably many path-components each of which consists of one end only;*
- *X has infinitely many path-components that contain a vertex; and*
- *every path-component of X contains an end.*

The counterexample can well be read by itself, but it may look somewhat surprising. However, the proofs of Theorems 4.3 and 4.1 will make it less surprising with hindsight: they will show why it had to be the way it is.

4.2 Connectedness does not imply path-connectedness in $|G|$

In this section we prove Theorem 4.2. Let $G = (V, E)$ be a graph. A subgraph consisting of a path xyz of order 3 and three disjoint rays starting at x, y, z respectively will be called a *trident*. The path xyz is the *spine* of the trident, and the rays are its *spikes*. The ends of G that contain the rays of the trident will be called, with slight abuse of terminology, the *ends* of the trident.

We will now recursively construct an infinite locally finite graph G and a subgraph X^* , which will be a collection of disjoint double rays of G , and will give rise to a connected but not path-connected subspace of $|G|$. At the same time we will define a sequence of trees $\{T_i\}_{i < \omega}$ of auxiliary use. All vertices of any T_i , apart from their common root r , will be tridents in G .

Start with two tridents t_0, t_1 with a common spine, but otherwise disjoint (Figure 4.2). Put the three disjoint double rays formed by their spikes in X^* . Let T_0 consist of its root r and t_0, t_1 each joined to r .




Figure 4.1: Three new tridents, with spikes μ, α, ν , and κ, β, λ , and o, γ, ξ

Now perform ω steps of the following type. At step i , consider separately every trident v in G that is a leaf of T_i . Denote the spikes of v as α, β, γ and add to G three disjoint double rays and 6 further edges as in Figure 4.1 (these 6 edges are shown in thin continuous lines) to obtain the three new tridents with spikes μ, α, ν , and κ, β, λ , and o, γ, ξ . Add these tridents to T_i as neighbours of v . Let T_{i+1} be the tree resulting from such addition of three new tridents at every leaf of T_i ; then T_{i+1} has no leaves in common with T_i . For every leaf of T_i , add to X^* the three double rays $\mu e_{\mu\xi}$, $\nu e_{\nu\kappa}$ and $\lambda e_{\lambda o}$ shown in dashed lines in Figure 4.1. (Note that the spikes α, β, γ of the old trident each contain a spike of one of the new tridents. Thus each ray will eventually participate in an infinite number of tridents.) Figure 4.2 shows the graph after the first and part of the second step.

Let G be the graph obtained after ω steps, let Ω denote its set of ends, and put $T = \bigcup_{n \in \mathbb{N}} T_n$. The vertices of T other than r will be called the *T-tridents*. We will call the countably many ends of G that contain some ray of a *T-trident* the *explicit* ends of G . Apart from them, G has continuum many other ends, which we will call *implicit*. They consist of rays that each meets infinitely many double rays of X^* .

We will construct a connected set $X \subset |G|$ that is not path-connected. The path-components of X containing vertices, demanded by Theorem 4.3, will be the closures of the double rays of X^* . In order to supply the singleton ends, we will now divide the implicit ends between X and its complement in




Figure 4.2: The first steps of the construction of G . The thick lines depict t_0 and t_1 .

$|G|$, in such a way that

neither X nor $\Omega \setminus X$ contains a closed set of continuum many ends. (*)

Since Ω has a countable basis (as a topological subspace of $|G|$), it has at most continuum many closed subsets. So we may index those closed subsets of Ω that contain continuum many ends as A_α , $\alpha < \gamma$ where γ is at most the initial ordinal of the continuum.

Then perform γ steps of the following type. At step α , use the fact that $|A_\alpha| \geq |\gamma| > |\alpha|$, and that only countably many ends in A_α are explicit, to pick two implicit ends from A_α that were not picked at any of the α earlier steps; earmark one of these ends for inclusion in X .

Define X as the union of all double rays in X^* , all explicit ends, and those implicit ends that have been earmarked. X clearly satisfies (*). We will show that X is a connected but not path-connected subspace of $|G|$, by proving the following implications:

- If X is not connected, then $\Omega \setminus X$ contains a closed set of continuum many ends.

- If X is path-connected, then X contains a closed set of continuum many ends.

In both cases, the validity of condition (*) is contradicted.

Let us prove the first implication. Suppose X is not connected; then X is contained in the union of two open sets O_r, O_g of $|G|$ which both meet X but whose intersection does not. Colour all points in $O_r \cap X$ red and all points in $O_g \cap X$ green. Note that every path-component of X , and in particular every double ray in X^* , is monochromatic, because it is a connected subspace of X .

If t is any T -trident with spine xyz and α one of its ends, then $U(t) := O(\{x, y, z\}, \alpha)$ is a basic open set that does not depend on the choice of α ; note that, by virtue of the ‘6 additional edges’ of Figure 4.1, all three spikes of t have a subray in the same component of $G - \{x, y, z\}$. Then $\mathcal{U} := \{U(t) | t \text{ is a } T\text{-trident}\}$ is a basis of the open neighbourhoods of the ends of G , because for every end and every finite $S \subset V(G)$ there is a $U(t)$ that contains the end and misses S .

Let us show that at least one of the T -tridents must contain vertices of both colours. If not, then all the vertices of t_0 and t_1 have the same colour, since double rays in X^* must be monochromatic. Moreover, every T -trident meets all its children in T , so all vertices of all T -tridents have the same colour, which means that $X^* \cap V$ is monochromatic. As \mathcal{U} is a basis of the open neighbourhoods of the ends, every open neighbourhood of an end meets $X^* \cap V$, so all ends in X (as well as, clearly, all edge points in X) also bear the colour of $X^* \cap V$, contradicting our assumption that both O_r, O_g meet X .

Next, we show that if a T -trident t is two-coloured, then there are two-coloured T -tridents r, s such that $U(r), U(s)$ are disjoint proper subsets of $U(t)$ (In other words r and s are both descendants of t , but not of each other). Let the tridents x, y, z be the children of t in T . We may assume that the spike of t that meets y is green, while its other two spikes are red (Figure 4.3).

Now consider the three thin double rays in Figure 4.3. If any of these is green, then at least two of the tridents x, y, z will be two-coloured. So let us assume that all those three double rays are red. But now y is coloured like t (one spike green, the other two red), and we may repeat the argument with y in the place of t . We continue recursively to find a descending ray $y_0 y_1 y_2 \dots$ in T (with $y_0 = t$ and $y_1 = y$) of two-coloured tridents. But the sets $U(y_i)$ form a neighbourhood basis of the end ω of the green spike of t . This contradicts the fact that $\omega \in O_g$ and O_g is open.

We have thus shown that T contains a subdivision B of the infinite binary




Figure 4.3: $U(t)$ and its subneighbourhoods.

tree all whose branch vertices are 2-coloured tridents. Let $\sigma = x_1x_2\dots$ be any descending sequence of branch vertices of B . Then $\bigcap_{i \in \mathbb{N}} U(x_i)$ contains a unique end, $\omega(\sigma)$. As the $U(x_i)$ form a neighbourhood basis of $\omega(\sigma)$ and are all 2-coloured, so $\sigma \in \Omega \setminus (O_r \cup O_g)$. Since B contains continuum many such sequences σ , and their corresponding ends $\omega(\sigma)$ are clearly distinct, the set Ω' of all these ends $\omega(\sigma)$ is a subset of $\Omega \setminus (O_r \cup O_g)$ containing continuum many ends. As $O_r \cup O_g$ is open, the closure of Ω' still lies in $\Omega \setminus (O_r \cup O_g) \subseteq \Omega \setminus X$. This contradicts (*), and completes our proof that X is connected.

It remains to prove that X is not path-connected. Suppose it is. Then any two distinct implicit ends $x, y \in X$ are connected by a path in X , and by Lemma 2.5 there is also an x - y arc A in X . We show that A contains continuum many ends, which will contradict (*).

It is easy to confirm (by Lemma 2.4) that A must contain a vertex of X^* . Clearly, the double ray $R \in X^*$ containing this vertex is a subarc of A . Let A' and A'' denote the path-components of $A \setminus R$, which are subarcs of A preceding and following R . As before, A' and A'' each contain a double ray from X^* , R' and R'' say. These double rays cannot share an end with R , because by construction no end contains more than one ray of X^* , hence R' and R'' split A' and A'' in two smaller subarcs.

Repeating recursively on each subarc of the previous step, we see that A contains a set \mathcal{R} of infinitely many double rays, arranged like the segments of the unit interval removed to form the Cantor set. Imitating the corresponding

proof, we see that A contains a set C of continuum many points that are limits of the ends of the double rays in \mathcal{R} . But only ends can be limits of ends, so C is a set of ends of X .

The arc A is closed because it is compact (as image of the compact space $[0, 1]$) and $|G|$ is a Hausdorff space (see [14] for a proof of this fact). The set of ends that lie on A is also closed, because its complement in $|G|$ consists of the complement of A in $|G|$ plus a set of vertex and edge points, and each of the later has an open neighbourhood that avoids all ends. Since this set contains C , it follows that A contains a closed subset of $\Omega \cap X$ with continuum many elements contradicting (*).

This completes the proof that X is not path-connected and hence the proof of Theorem 4.2.

4.3 Connectedness implies path-connectedness almost always

In this section, X will denote an arbitrary connected subspace of $|G|$, where $G = (V, E)$ is an arbitrary locally finite connected graph. We assume that X does not entirely lie on an edge of G , in which case it would obviously be path-connected.

The aim of this section is to prove Theorem 4.3. To this end we will first have to develop some intermediate results.

For $x \in X$, let $c(x)$ denote the path-component of X that contains x .

Lemma 4.1. *For every point $x \in X \setminus \Omega$ there is an open neighbourhood $U = U(x)$ of x such that $U \cap X \subseteq c(x)$.*

Proof. First assume that x is an inner point of the edge $[u, v]$. We claim that

$$\text{one of the closed intervals } [u, x], [x, v] \text{ lies in } X \text{ as well.} \quad (4.1)$$

For suppose not. Then there is a point $u' \in [u, x)$ and a point $v' \in (x, v]$ that do not belong to X . But then (u', v') and $|G| \setminus [u', v']$ are disjoint open subsets of $|G|$ that both meet X and whose union contains X , contradicting the connectedness of X .

Thus (4.1) holds and we may assume without loss of generality that $[u, x] \subset X$. Now if X contains an interval $(x, w) \subset [x, v]$ we can set $U(x) = (u, w)$. Otherwise there is a point $v' \in (x, v)$ such that $(x, v') \cap X = \emptyset$, and we can set $U(x) = (u, v')$. For if no such v' exists, then there are points of X on (x, v) arbitrarily close to x . But for every such point y we can prove that $[y, v] \subset X$ the same way we proved (4.1) ($[u, y] \not\subset X$ because

then we would have the previous case) and thus $[x, v] \subset X$ contradicting the assumption that X contains no interval $(x, w) \subset [x, v]$.

Now assume x is a vertex. By a similar argument as above we can prove that for every edge $xv \in E(G)$ there is a point $v' \in (x, v)$ such that either $(x, v') \subset X$ or $(x, v') \cap X = \emptyset$. Let S be a set that contains one such point for each edge incident with x . Then we can set $U(x) = \bigcup_{v' \in S} [x, v')$. \square

For an end $\omega \in X$ we cannot in general find a neighbourhood of ω that meets X only in $c(\omega)$. However, we can always find one that avoids any specified path-component other than $c(\omega)$:

Lemma 4.2. *For every end $\omega \in X$ and every path-component $c' \neq c(\omega)$ of X there is an open neighbourhood $U = U(c', \omega)$ of ω such that $U \cap c' = \emptyset$.*

In order to prove this lemma, we will suppose that there is a path-component c' of X and an end $\omega \in X \setminus c'$ every neighbourhood of which meets c' . To construct a path from c' to ω in X contradicting $c' \neq c(\omega)$, we shall pick a sequence a_0, a_1, a_2, \dots of vertices in c' converging to ω , link a_i to a_{i+1} by a path in c' for each i , and concatenate all these paths to a map $f : [0, 1) \rightarrow c'$. Adding $f(1) := \omega$ yields an a_0 - ω path in X as long as f is continuous at 1. To ensure this, we have to choose our a_i - a_{i+1} paths inside smaller and smaller neighbourhoods U_i of ω .

Proof of Lemma 4.2. Suppose there is a path-component c' of X and an end $\omega \in X \setminus c'$ every open neighbourhood of which meets c' . Easily by Lemma 2.4, c' must contain a vertex u .

Define $S_0 = \emptyset$, and for every $i > 0$ let $S_i = N^{i-1}(u)$. Let $U_i = O(S_i, \omega)$. Note that $S_0 \subset S_1 \subset S_2 \subset \dots$, and thus $U_0 \supset U_1 \supset U_2 \supset \dots$.

Define $M_i = (S_{i+1} \setminus S_i) \cap c' \cap U_i$, for all $i \geq 0$ (Figure 4.4). Each M_i is a set of candidates for the vertex a_i mentioned above. Instead of choosing them arbitrarily, we will make use of Lemma 2.7 to find a sequence of appropriate a_i .

Define the graph \mathcal{G} with $V(\mathcal{G}) = \bigcup_i M_i$ and $xy \in E(\mathcal{G})$ if for some i , $x \in M_i, y \in M_{i-1}$ and there is a x - y topological path in $c' \cap U_{i-1}$.

We need to show that \mathcal{G} satisfies the conditions of Lemma 2.7. Since G is locally finite, the S_i are finite, and hence so are the M_i . Let us show that they are non-empty.

For $i > 0$ pick any point of $U_i \cap c'$ and any topological path from that point to u . By Lemma 2.4, and since $u \notin U_i$, this path traverses one of the edges between a vertex w in U_i and a vertex outside it. By definition, M_i contains this vertex w .

In order to see that every $x \in M_i$ sends an edge to M_{i-1} , pick any $z \in M_{i-1}$, and any topological path in c' from x to z . Since M_{i-1} is closed,




Figure 4.4: U_i and M_i .

this path has a first point y in M_{i-1} . By Lemma 2.4, the subpath from x to y lies in U_{i-1} , so xy is an edge of \mathcal{G} .

We can now apply Lemma 2.7 to get an infinite path $a_0(=u)a_1a_2\dots$ in \mathcal{G} . For each $i > 0$, pick a topological path f_i in $c' \cap U_{i-1}$ from a_{i-1} to a_i (which exists because $a_{i-1}a_i$ is an edge of \mathcal{G}), let $f : [0, 1) \rightarrow c'$ be the concatenation of these paths, and put $f(1) = \omega$.

We claim that f is continuous at 1 and hence a path in X , contradicting our assumption that $c' \neq c(\omega)$. To see that this is the case, let O be any open neighbourhood of ω . Choose a basic open neighbourhood $O' = O(S, \omega) \subseteq O$. Let i be the maximum distance of an element of S from u . Then $S_i \supseteq S$, and $U_i \subseteq O' \subseteq O$. Since for $j > i$ the path f_j lies in $U_{j-1} \subseteq U_i$, the subpath of f from a_i to $f(1)$ lies in $U_i \subseteq O$ which proves the continuity of f at 1. This completes the proof. \square

As a consequence of Lemmas 4.1 and 4.2 we have the following:

Lemma 4.3. *The path-components of X are closed in its subspace topology.*

This implies that any counterexample to Problem 4.1 must contain infinitely many path-components. In fact we can prove something stronger:

Lemma 4.4. *Every connected but not path-connected $X \subseteq |G|$ contains uncountably many path-components.*

Proof. Suppose c_1, c_2, \dots is an enumeration of the path-components of X . We will divide X into two open sets O_r, O_g of $|G|$ whose intersection does not meet X contradicting its connectedness.

We will proceed recursively. Every path-component c will at some step be coloured either red or green (Eventually, O_r will be a union of open sets that

contains all points that belong to red path-components, and O_g similarly for ‘green’). If c is not immediately put in one of O_r, O_g (as part of some open set) at the step that it gets coloured, it will be given a natural number as *handicap*. This handicap will be a competitive advantage for the ends in c against ends whose path-component has a higher handicap, and which are also striving to get classified in O_r or O_g , and will help make sure that every end in c will be classified after a finite number of steps (but if c has infinitely many ends, it might take infinitely many steps till they all get classified).

Once we have accommodated all ends of X in either of O_r, O_g it will be easy to do the same for the vertices and edge points of X .

At the beginning of step i of the recursion we will pick a finite set $S_i \subset V$, which grows larger at each step, and consider the (finitely many) open sets of the form $O(S_i, \omega)$, for all $\omega \in \Omega$. We will declare *live* any such open set that contains ends of X that have not yet been classified in O_r or O_g . Some of these open sets might be put in O_r or O_g during the current step, in which case we will switch their state to not live. Each live open set L will have a *boss*, namely, the path-component of smallest handicap meeting L . Being a boss will let a path-component influence subsequent colouring decisions for its own ends.

Formally, we apply the following recursion. Before the first step, colour c_0 red and c_1 green; this will guarantee that neither of O_r, O_g will be empty. Give c_0 the handicap 0, and c_1 the handicap 1. Declare $|G|$ live, and let c_0 be its boss. Let u be any vertex of G .

Then for every $i < \omega$ perform the following actions (see Figures 4.5 and 4.6):

1. Declare live all the basic open sets of the form $O(N^i(u), \omega)$, with $\omega \in \Omega \cap X$ that lie in live open sets of the previous step (note that $O(N^i(u), \omega) \subsetneq O(N^{i-1}(u), \omega)$).
2. Colour any still uncoloured path-component c that meets more than one live open set with the colour of the boss of the parent open set, i.e. the live open set of the previous step in which c lies (it must lie in one, because if it met more than one of them it would have been coloured in a previous step). Note that there are only finitely many such path-components in any step, because by Lemma 2.4 each of them must contain an edge that crosses some basic open set and there are only finitely many such edges. Finally give the newly coloured path-components the next free handicaps, one to each.
3. If a live basic open set does not meet any green path-components, then colour all path-components that lie in it red, put it in O_r and declare

it not live. Proceed similarly with colours switched and O_g instead of O_r .

4. For every live basic open set, let the path-component of smallest handicap that meets it be its boss.
5. If c_i is still uncoloured, give it the colour of the boss of the live set in which it lies (it lies in one since it is still uncoloured) and the next free handicap.




Figure 4.5: Possible colourings after step 1. Dashed lines depict red path-components and continuous lines green ones. The path-component c_k meets several live sets, so it took the colour of c_0 , the boss of $|G|$. The basic open set C_1 will be put in O_r and C_2 will be put in O_g ; then they will be declared not live. The boss of both C_3 and C_4 is c_1 .

We claim that after this process every end of X is put in either O_r or O_g . Indeed, because of action 5, for every end e of X , $c(e)$ gets a colour and a handicap sometime. By Lemma 4.2 and the fact that there are only finitely many path-components of smaller handicap, at some step j , e will lie in a live basic open set U that avoids all path-components of smaller handicap ($N^i(u)$ contains any finite vertex set for i large enough, if we assume, without loss of generality, that G is connected).




Figure 4.6: Possible colourings after step 2. $C_{4,1}$ will be put in O_r and $C_{3,2}$ in O_g . The path-component c_2 received the colour of c_1 , the boss of $C_{4,2}$, and c_l received the colour of the boss of C_3 , again c_1 . The arrows show the bosses of the open sets that are live after the completion of this step.

At this point, U only meets finitely many coloured path-components (see comment in action 2). In the steps following step j , e 's path-component will always be the boss of the current live open set in which e lies (action 4) and thus no path-component that meets such a set will be coloured with the opposite colour (action 2).

Again by Lemma 4.2, e will at some later step lie in a basic open set U' that avoids all path-components of the opposite colour that met U at step j . This U' thus meets only the colour of e , so it will be classified in one of O_r, O_g .

Thus our claim is true and we have divided $X \cap \Omega$ into two open sets whose intersection does not meet X . Now for each vertex or edge point of X find a basic open set, supplied by Lemma 4.1, that avoids all other path-components, and put it in O_r if the point belongs to a red path-component, or in O_g if it belongs to a green one. Since $X \subseteq O_r \cup O_g$, and $O_r \cap O_g$ does indeed not meet X , the connectedness of X is contradicted.

□

Using the above Lemmas we can now prove Theorem 4.3:

Proof of Theorem 4.3. Suppose X is not path-connected. Since G is locally finite and connected, X has only countably many path-components containing vertices, so by Lemmas 4.4 and 2.4 there must be uncountably many path-components that are singleton ends.

If c_1, c_2, \dots, c_n are the only path-components of X that contain vertices, then pick a singleton end $\omega \in X$, and for each i an open neighbourhood O_i of ω that avoids c_i , supplied by Lemma 4.2. Let $O = O(S, \omega)$ be a basic open neighbourhood of ω contained in $\bigcap_{i \leq n} O_i$ (if $n = 0$ then let O be any basic open neighbourhood of ω that avoids at least one end of X). Every point of $X \setminus O$ has an open neighbourhood that does not meet O : for vertices and edge points this open neighbourhood is supplied by Lemma 4.1 and for an end $e \notin O$ the neighbourhood $O = O(S, e)$ does indeed not meet O . Thus $X \cap O$ is open and closed in the subspace topology of X , a contradiction since X is connected. This proves that X must have infinitely many path-components that contain vertices.

Finally, let us show that every path-component of X contains ends. By Lemma 4.1, a path-component containing no end is open. Since it is also closed (Lemma 4.3), the connectedness of X is contradicted if such a path-component exists.

□

4.4 Connectedness implies path-connectedness for closed subspaces

In this section we give a proof of Theorem 4.1. As already mentioned, Theorem 4.1 was already proved in [16]¹. Because of its importance for applications, we present a simpler proof that makes use of the results of Section 4.3. This proof could considerably facilitate the understanding of why the counterexample in Section 4.2 has to be the way it is.

We will need some definitions and a lemma. If S is a finite subset of $V(G)$, we will say that two ends $x, y \in X \cap \Omega$ are *S-equivalent*, and write $x \sim_S y$, if there is a sequence $(x =) \omega_1 \omega_2 \dots \omega_k (= y)$, called a *x-y connecting sequence* (relative to S), such that for each i , either there is a double ray in X with ends ω_i, ω_{i+1} , or ω_i, ω_{i+1} lie in the same (topological) component of $X \cap O(S, \omega_i)$ (in particular, $O(S, \omega_{i+1}) = O(S, \omega_i)$). We will say that a path-component *crosses* a set $O \subset |G|$, if it meets both O and $|G| - O$.

¹In fact the result in [16] is more general, as it applies not only to locally finite graphs but also to a class of non-locally-finite ones.

Lemma 4.5. *For every finite $S \subset V(G)$, every two ends x, y in X are S -equivalent.*

Proof. Suppose for contradiction that x, y are not S -equivalent. Clearly, \sim_S is an equivalence relation, so let C_1 be the equivalence class of x and let $C_2 = X \cap \Omega - C_1$.

We will split X into disjoint open sets O_1, O_2 , with $x \in O_1$ and $y \in O_2$. Let \mathcal{O} be the (finite) set of open sets $O(S, \omega), \omega \in \Omega$. For every $Q \in \mathcal{O}$, there are only finitely many path-components of X crossing Q , because S is finite. This and the definition of \sim_S implies that there are only finitely many (topological) components of $Q \cap X$ meeting C_1 . So if Q meets both C_1, C_2 , then we can find non-empty disjoint open subsets $Q_1 = Q_1(Q), Q_2 = Q_2(Q)$ of $Q \cap X$ so that $Q \cap X = Q_1 \cup Q_2, C_1 \cap Q \subseteq Q_1$, and no end in $C_2 \cap Q_1$ belongs to a path-component of X crossing Q . If Q does not meet C_1 (respectively C_2), let $Q_1 = \emptyset, Q_2 = Q \cap X$ (resp. $Q_1 = Q \cap X, Q_2 = \emptyset$).

Let R_1 (respectively R_2) be the set of vertices and edge points in $X - \bigcup \mathcal{O}$ lying in a path component of X that meets C_1 (resp. C_2). By Theorem 4.3, every point of $X - \Omega$ is connected to an end of X by a path, so $R_1 \cup R_2 = X - \bigcup \mathcal{O}$. Moreover, no path-component of X can meet both C_1, C_2 , as a path between two ends x, y induces an x - y connecting sequence for every S , thus $R_1 \cap R_2 = \emptyset$ (here we used the fact that a path-connected topological space is connected).

Now let $O_1 = \bigcup_{Q \in \mathcal{O}} Q_1(Q) \cup \bigcup_{r \in R_1} U(r)$, where $U(r)$ is an open neighbourhood of r such that $U(r) \cap X \subseteq c(r)$, supplied by Lemma 4.1. Similarly, let $O_2 = \bigcup_{Q \in \mathcal{O}} Q_2(Q) \cup \bigcup_{r \in R_2} U(r)$. By construction, O_1, O_2 are open, non-empty, disjoint and $O_1 \cup O_2 \supseteq X$, contradicting the connectedness of X . \square

We can now proceed to our proof.

Proof of Theorem 4.1. Let $x, y \in X \cap \Omega$. We will construct an x - y topological path σ in X . Since by Theorem 4.3, every path-component of X contains an end, the assertion follows.

Pick an arbitrary vertex $u \in V(G)$, and for every $n \in \mathbb{N}$ define $S_n = G[u]^n$. Applying Lemma 4.5 for $S = S_0$, we obtain a x - y connecting sequence $(x =)\omega_0\omega_1 \dots \omega_k(= y)$. As a first approach to defining σ , let $\sigma_0 : [0, 1] \rightarrow X \cup \{\emptyset\}$ be a mapping, such that $\sigma(\frac{i}{k}) = \omega_i$ for every $0 \leq i \leq k$, and if there is a double ray R with ends ω_i, ω_{i+1} , then σ maps $[\frac{i}{k}, \frac{i+1}{k}]$ continuously onto R . If there is no double ray with ends ω_i, ω_{i+1} , then σ maps $(\frac{i}{k}, \frac{i+1}{k})$ to \emptyset .

Now recursively, for $n = 1, 2, \dots$, suppose that σ_{n-1} has already been defined, and define the mapping $\sigma_n : [0, 1] \rightarrow X \cup \{\emptyset\}$ as follows. For every $p \in [0, 1]$, if $\sigma_{n-1}(p) \neq \emptyset$ then let $\sigma_n(p) = \sigma_{n-1}(x)$. For every maximal interval $I = (p, q)$ of $[0, 1]$ mapped to \emptyset , it follows by construction that

$p, q \in X \cap \Omega$, and that p, q lie in the same (topological) component X' of $X \cap O(S_{n-1}, p)$. Now apply Lemma 4.5 for $S = S_n$ replacing X by X' , to obtain a p - q connecting sequence relative to S_n , where all double rays lie in $O(S_{n-1}, p)$. Then, let σ_n map I to double rays between ends of this connecting sequence and to \emptyset (if needed), following the pattern according to which σ_0 was defined.

Note that if $\sigma_i(p) \neq \emptyset$ for some $p \in [0, 1]$, then $\sigma_n(p) = \sigma_i(p)$ holds for every $n \geq i$. Thus we can define a mapping $\sigma' : [0, 1] \rightarrow X \cup \{\emptyset\}$, by letting $\sigma'(p) = \emptyset$ if $\sigma_n(p) = \emptyset$ for every n , and $\sigma'(p) = x$ if $\sigma_n(p) = x \neq \emptyset$ for some n . Since any two ends in Ω are separated by some S_n , no open interval of $[0, 1]$ is mapped to \emptyset by σ' . Moreover, for every point p such that $\sigma'(p) = \emptyset$, there is a descending sequence a of open sets of the form $O(S_n, \omega)$, such that an open interval around p is mapped to each member of a . As $|G|$ is compact, a has an accumulation point, and since any two ends are separated by some S_n , a converges to an end $\omega_p \in \Omega$. But X is closed, so $\omega_p \in X$. Now define the mapping $\sigma : [0, 1] \rightarrow X$, by letting $\sigma(p) = \omega_p$ if $\sigma'(p) = \emptyset$ and $\sigma(p) = \sigma'(p)$ otherwise. By construction, σ is an x - y topological path in X , which completes the proof.

□

Chapter 5

Linear algebra with infinite sums

5.1 Introduction

In this chapter, which is based on [5], we will answer the following questions: given a set \mathcal{N} of edge sets of some infinite graph, for instance a set of circuits, does \mathcal{N} have a basis $\mathcal{B} \subseteq \mathcal{N}$, in the sense that any edge set that is a sum of a thin family of elements of \mathcal{N} , can also be written as a sum of a thin family of elements of \mathcal{B} , and \mathcal{B} is minimal with that property? If $\langle \mathcal{N} \rangle$ is the set of all sums of thin families of elements of \mathcal{N} , then is $\langle \mathcal{N} \rangle$ closed under taking thin sums?

These problems have applications in graph theory, as we shall see, but are also interesting from the algebraic point of view, so rather than confining ourselves to edge sets, we will in fact let \mathcal{N} be a set of subsets of any abstract set M in the above questions.

An element N of $\mathcal{P}(M)$ can also be represented as an element $\Phi(N)$ of \mathbb{Z}_2^M : just let $\Phi(N)(m) = 1$ if $m \in N$ and $\Phi(N)(m) = 0$ otherwise. With this observation, our setting becomes reminiscent of linear algebra. We have a “vector space” over \mathbb{Z}_2 , where some sums of infinitely many summands are allowed. In order to make our treatment more general, we will replace \mathbb{Z}_2 by an arbitrary field or ring R .

After formally defining the concepts in use, we will show that the answer to our first question is positive if R is a field and M is countable (Theorem 5.1), but not otherwise (Theorem 5.2). The answer to our second question is easily seen to be negative if \mathcal{N} is not thin, that is, it contains infinitely many elements that “meet” a certain $m \in M$. If however \mathcal{N} is thin, then the answer is positive if R is a field (Theorem 5.3), or if R is a finite ring

(Theorem 5.4), but not otherwise (Theorem 5.5).

Theorems 5.1 and 5.3 can be useful tools in infinite graph theory, in particular in the study of the topological cycle space $\mathcal{C}(G)$. For example, the existence of bases (in our sense) was instrumental in [6], where MacLane’s planarity criterion was generalised to locally finite graphs. Theorem 5.3 provides an alternative — and easier — way of proving that $\mathcal{C}(G)$ is closed under taking thin sums, which was first shown in [14, Corollary 5.2]: indeed, it suffices to apply Theorem 5.3 on the thin set \mathcal{N} of fundamental circuits of a normal spanning tree, which by [12, Lemma 8.5.7] and [12, Theorem 8.5.8] generate $\mathcal{C}(G)$.

5.2 Bases

Let M be a set, R be a ring, and let $\mathcal{L} \subseteq R^M$. We call a function $a : \mathcal{L} \rightarrow R$ *thin* if for every $m \in M$, there are only finitely many $N \in \mathcal{L}$ such that $a(N)N(m) \neq 0$. Thin functions are precisely those functions $a : \mathcal{L} \rightarrow R$ for which the *sum* $\sum a := \sum_{N \in \mathcal{L}} a(N)N$ of a is well defined. For a $K \in R^M$, we call a (thin) function $a : \mathcal{L} \rightarrow R$ a *representation of K in \mathcal{L}* if $K = \sum a$ — that is, $K(m) = \sum_{N \in \mathcal{L}} a(N)N(m)$ for every $m \in M$. Denote by $\langle \mathcal{L} \rangle$ the set of elements of R^M that have a representation in \mathcal{L} . Equivalently, $\langle \mathcal{L} \rangle$ is the set of sums of thin functions $a : \mathcal{L} \rightarrow R$. Intuitively, \mathcal{L} is a generating set, and $\langle \mathcal{L} \rangle$ is the space it generates.

For a set $\mathcal{N} \subseteq R^M$, we call a subset \mathcal{B} of \mathcal{N} a *basis of $\langle \mathcal{N} \rangle$* , if $\langle \mathcal{B} \rangle = \langle \mathcal{N} \rangle$ and \mathcal{B} is minimal with that property. If R is a field, then this is equivalent to saying that $\langle \mathcal{B} \rangle = \langle \mathcal{N} \rangle$ and $\mathbf{0}$, the zero function in R^M , has a unique representation in \mathcal{B} (namely the function mapping every $N \in \mathcal{N}$ to $0 \in R$). Thus, if \mathcal{B} is a basis of \mathcal{N} and R is a field, then every element of $\langle \mathcal{N} \rangle$ has a unique representation in \mathcal{B} .

Our first result states that if R is a field and M is countable, then we can always find a basis:

Theorem 5.1 ([5]). *Let M be a countable set, F be a field and let $\mathcal{N} \subseteq F^M$. Then \mathcal{N} contains a basis of $\langle \mathcal{N} \rangle$.*

In linear algebra the analogous assertion is usually proved with Zorn’s lemma as follows. Given a chain $(\mathcal{B}_\lambda)_\lambda$ of linearly independent subsets of the generating set, it is observed that $\bigcup_\lambda \mathcal{B}_\lambda$ is still linearly independent since any violation of linear independence is witnessed by finitely many elements, and these would already lie in one of the \mathcal{B}_λ . Thus, each chain has an upper bound, which implies, by Zorn’s lemma, that there is a maximal linearly independent set, a basis. This approach however fails in our context, as

dependence does not need to be witnessed by only finitely many elements, thus we cannot get the contradiction that already one of the \mathcal{B}_λ was not independent.

As an illustration, put $F = \mathbb{Z}_2$, $M := \mathbb{Z}$ and $\mathcal{B}_i := \{\Phi(\{j, j+1\}) : -i \leq j < i\}$ for $i = 1, 2, \dots$. Now, while no nonempty finite subset of $\mathcal{B}_\infty := \bigcup_{i=1}^\infty \mathcal{B}_i$ is dependent, the whole set is: $\sum_{B \in \mathcal{B}_\infty} B = \mathbf{0}$.

Proof of Theorem 5.1. If M is a finite set the result follows from linear algebra as no infinite sum is well defined, so we assume M to be infinite. Let m_1, m_2, \dots be an enumeration of M , and for $i = 1, 2, \dots$ define \mathcal{N}_i to be the set of those elements $N \in \mathcal{N} \setminus \bigcup_{j < i} \mathcal{N}_j$ for which $N(m_i) \neq 0$. Clearly, $\{\mathcal{N}_i : i \in \mathbb{N}\}$ is a partition of \mathcal{N} . For every $i \in \mathbb{N}$, let $N_{i1}, N_{i2}, \dots, N_{i\lambda}, \dots$ be a (possibly transfinite) enumeration of \mathcal{N}_i .

Let us briefly outline the proof. For each $i \in \mathbb{N}$ we will inductively go through the elements of \mathcal{N}_i , and each time we encounter a $N_{i\lambda}$ that has a representation that only uses (i.e. takes a nonzero value on) the predecessors of $N_{i\lambda}$ in \mathcal{N}_i (that is, the $N_{i\mu}$ for which $\mu < \lambda$) and elements of later \mathcal{N}_j , we will delete $N_{i\lambda}$. Doing this for every i ensures that $\mathbf{0}$ has a unique representation in the remaining subset of \mathcal{N} .

We then have to check that we can still represent any element K of $\langle \mathcal{N} \rangle$. This works as follows. Let a be a representation of K in \mathcal{N} , and let i be the least integer so that a uses some deleted elements of \mathcal{N}_i . Each time we deleted a $N_{i\lambda}$, we made sure that the elements of \mathcal{N}_i needed to represent it would not be deleted afterwards (because they had smaller indices). We could, of course, delete elements of \mathcal{N}_{i+1} needed to represent $N_{i\lambda}$, but when we do, we make sure that we have left enough elements of \mathcal{N}_{i+1} behind, as we did with \mathcal{N}_i , so that at the end the remaining elements of all the \mathcal{N}_j can represent $N_{i\lambda}$.

Formally, perform ω steps of the following type. In step i , perform the following transfinite recursion. Start by setting $\mathcal{N}_{i0} = \mathcal{N}_i$, and then for every ordinal $\lambda > 0$ define the set $\mathcal{N}_{i\lambda} \subseteq \mathcal{N}_i$ as follows (intuitively, $\mathcal{N}_{i\lambda}$ is the set of elements of \mathcal{N}_i that we have not deleted so far): If $N = N_{i\lambda}$ has a representation in

$$X_{i\lambda} := \left(\bigcap_{\mu < \lambda} \mathcal{N}_{i\mu} \setminus \{N_{i\mu} : \mu \geq \lambda\} \right) \cup \bigcup_{k > i} \mathcal{N}_k$$

(that is, if N has a representation in $\bigcup_{k \geq i} \mathcal{N}_k$ that does not use any elements of \mathcal{N}_i that have already been deleted or have index greater than λ), then let $\mathcal{N}_{i\lambda} := \bigcap_{\mu < \lambda} \mathcal{N}_{i\mu} \setminus \{N_{i\lambda}\}$ (this corresponds to deleting $N_{i\lambda}$), and let a_N be a representation of N in \mathcal{N} such that

$$a_N(L) = 0 \text{ if } L \notin X_{i\lambda}, \quad (5.1)$$

which exists since N has a representation in $X_{i\lambda}$.

Otherwise, let $\mathcal{N}_{i\lambda} := \bigcap_{\mu < \lambda} \mathcal{N}_{i\mu}$. Having defined all $\mathcal{N}_{i\lambda}$, we put $\mathcal{B}_i := \bigcap_{\lambda} \mathcal{N}_{i\lambda}$.

We claim that $\mathcal{B} := \bigcup_i \mathcal{B}_i$ is a basis of $\langle \mathcal{N} \rangle$. To show that $\mathbf{0}$ has a unique representation, suppose there is a nonzero thin function $b : \mathcal{B} \rightarrow F$ such that $\sum_{C \in \mathcal{B}} b(C)C = 0$. Let $i \in \mathbb{N}$ be minimal so that there is an ordinal μ with $b(N_{i\mu}) \neq 0$, and observe that since for all the elements B in \mathcal{B}_i we have $B(m_i) \neq 0$, there is a maximal ordinal λ such that $b(N_{i\lambda}) \neq 0$. Then $\{N \in \mathcal{N} : b(N) \neq 0\} \subseteq X_{i\lambda}$, a contradiction to that $N_{i\lambda} \in \mathcal{B}_i \subseteq \mathcal{N}_{i\lambda}$.

Next, consider a $K \in \langle \mathcal{N} \rangle$. We will show that K has a representation in \mathcal{B} . Let b^0 be a representation of K in $\langle \mathcal{N} \rangle$. Inductively, define for $k = 1, 2, \dots$ thin functions $b^k : \mathcal{N} \rightarrow F$ as follows. (Intuitively, b^k is a representation of K using only elements of \mathcal{N} that are left after step k of the construction of \mathcal{B} , that is, after we have finished deleting elements of \mathcal{N}_k .) Let $\mathcal{E}_k = \{N \in \mathcal{N}_k \setminus \mathcal{B}_k : b^{k-1}(N) \neq 0\}$. Since b^{k-1} is thin and since $N(m_k) \neq 0$ for all $N \in \mathcal{E}_k \subseteq \mathcal{N}_k$, it follows that \mathcal{E}_k is a finite set. Put

$$b^k(N) = 0 \text{ for } N \in \mathcal{E}_k, \text{ and} \quad (5.2)$$

$$b^k(N) = b^{k-1}(N) + \sum_{L \in \mathcal{E}_k} b^{k-1}(L)a_L(N) \text{ for } N \notin \mathcal{E}_k. \quad (5.3)$$

(Note that a_L is well defined for every $L \in \mathcal{E}_k$, as $\mathcal{E}_k \subset \mathcal{N}_k \setminus \mathcal{B}_k$.)

By induction we easily obtain that for every $k \geq 1$

$$b^k(N) = 0 \text{ if } N \in \bigcup_{l=1}^k \mathcal{N}_l \setminus \mathcal{B} \quad (5.4)$$

and

$$K = \sum_{N \in \mathcal{N}} b^k(N)N \text{ (in particular, } b^k \text{ is thin)}. \quad (5.5)$$

Indeed, for (5.4) let $N = N_{l\lambda} \in \mathcal{N}_l \setminus \mathcal{B}$ with $1 \leq l \leq k$. If $b^{l-1}(N) \neq 0$, then $N \in \mathcal{E}_l$ and thus $b^l(N) = 0$ by (5.2). If $b^{l-1}(N) = 0$, then by (5.1) we have $a_L(N) = 0$ for every $L \in \mathcal{E}_l \subseteq \mathcal{N}_l$; indeed, for any such $L = N_{l\nu}$, if $N \notin \{N_{l\mu} : \mu \geq \nu\}$, then $N \notin \bigcap_{\mu < \lambda} \mathcal{N}_{l\mu}$, as N was deleted, i.e. $N \notin \mathcal{N}_{l\lambda}$, so $N \notin X_{l\nu}$. Thus by (5.3) $b^l(N) = 0$ again. Since $a_L(N) = 0$ for every $L \in \mathcal{E}_i \subseteq \mathcal{N}_i$ with $i > l$ by (5.1), we obtain inductively by (5.3) that $b^k(N) = 0$ for every $k \geq l$.

For (5.5) we have

$$\begin{aligned}
& \sum b^k - \sum b^{k-1} = \sum_{N \in \mathcal{N}} (b^k(N)N - b^{k-1}(N)N) \\
= & \sum_{N \in \mathcal{N} \setminus \mathcal{E}_k} (b^k(N)N - b^{k-1}(N)N) + \sum_{N \in \mathcal{E}_k} (b^k(N)N - b^{k-1}(N)N) \quad (5.6) \\
= & \sum_{N \in \mathcal{N} \setminus \mathcal{E}_k} \left(\sum_{L \in \mathcal{E}_k} b^{k-1}(L) a_L(N) N \right) + \sum_{N \in \mathcal{E}_k} (\mathbf{0} - b^{k-1}(N)N).
\end{aligned}$$

As $a_L(N) = 0$ if $N, L \in \mathcal{E}_k$ by (5.1), we have

$$\begin{aligned}
\sum_{N \in \mathcal{N} \setminus \mathcal{E}_k} \left(\sum_{L \in \mathcal{E}_k} b^{k-1}(L) a_L(N) N \right) &= \sum_{L \in \mathcal{E}_k} b^{k-1}(L) \sum_{N \in \mathcal{N} \setminus \mathcal{E}_k} a_L(N) N \quad (5.7) \\
&= \sum_{L \in \mathcal{E}_k} b^{k-1}(L) \sum_{N \in \mathcal{N}} a_L(N) N = \sum_{L \in \mathcal{E}_k} b^{k-1}(L) L.
\end{aligned}$$

By (5.6) and (5.7) we get $\sum b^k - \sum b^{k-1} = \mathbf{0}$, which proves (5.5).

For every $N \in \mathcal{N}_l$ define $b^\infty(N) := b^l(N)$, and note that

$$b^k(N) = b^\infty(N) \text{ for } N \in \mathcal{N}_l \text{ and } k \geq l. \quad (5.8)$$

Indeed, consider $k > l$ and observe that, by (5.1), $a_L(N) = 0$ for all $L \in \mathcal{E}_k$, so by (5.3) we get $b^k(N) = b^{k-1}(N)$.

From (5.8) and (5.4) we easily get for $N \in \mathcal{N}$ that

$$\text{if } N \notin \mathcal{B} \text{ then } b^\infty(N) = 0. \quad (5.9)$$

We claim that

$$K = \sum_{B \in \mathcal{B}} b^\infty(B)B \text{ (in particular, } b^\infty \text{ is thin)}. \quad (5.10)$$

Consider an $m_k \in M$. As $N(m_k) = 0$ for $N \in \mathcal{N}_l$ with $l > k$, we get from (5.8) that $b^\infty(N)N(m_k) = b^k(N)N(m_k)$ for all $N \in \mathcal{N}$. Thus by (5.5) we obtain

$$\sum_{N \in \mathcal{N}} b^\infty(N)N(m_k) = \sum_{N \in \mathcal{N}} b^k(N)N(m_k) = K(m_k).$$

Claim (5.10) now follows from (5.9). This completes the proof. \square

Observe that contrary to traditional linear algebra, two bases do not need to have the same cardinality. Indeed, putting $F = \mathbb{Z}_2$ and $M = \{m_0, m_1, \dots\}$ we see that $\mathcal{B} := \{\Phi(\{m_i\}) : i \geq 0\}$ is a countable basis of F^M . On the other hand, $\mathcal{N} := \{\Phi(\{m_0\} \cup N) : N \subseteq M\}$ clearly generates F^M , and contains, by Theorem 5.1, a basis \mathcal{B}' . Since all thin subsets of \mathcal{N} are finite, \mathcal{B}' needs to be uncountable to generate the uncountable set F^M . Thus \mathcal{B} and \mathcal{B}' are two bases of F^M that do not have the same cardinality.

We have formulated Theorem 5.1 only for countable sets M . The following result shows that this is indeed best possible.

Theorem 5.2 ([5]). *There is an uncountable set M and $\mathcal{N} \subseteq \mathbb{Z}_2^M$ so that \mathcal{N} does not contain a basis of $\langle \mathcal{N} \rangle$.*

Proof. In order to simplify our expressions in this proof, we want to work with $\mathcal{P}(M)$ rather than \mathbb{Z}_2^M . As Φ is a bijection between the two sets respecting addition and thinness, we may do so. Thus we will apply expressions defined for (sets of) elements of \mathbb{Z}_2^M on (sets of) elements of $\mathcal{P}(M)$. These expressions inherit then their meaning from our previous definitions, if we “translate” elements of $\mathcal{P}(M)$ into elements of \mathbb{Z}_2^M .

Let A, B be two disjoint sets with cardinalities $|A| = \aleph_0$ and $|B| = \aleph_1$. Define G to be the graph with vertex set $M := A \cup B$ and edge set $\mathcal{N} := A \times B$. As $\mathcal{N} \subseteq \mathcal{P}(M)$, we may ask whether \mathcal{N} contains a basis of $\langle \mathcal{N} \rangle$ (that is, whether $\mathcal{N}' := \{\Phi(N) : N \in \mathcal{N}\}$ contains a basis of $\langle \mathcal{N}' \rangle$). We claim that it does not.

Let us show that each countable subset N of M is contained in $\langle \mathcal{N} \rangle$. Indeed, let n_1, n_2, \dots be a (possibly finite) enumeration of N , and choose for $i = 1, 2, \dots$ a ray R_i that starts at n_i , does not meet the first $i - 1$ vertices of each of R_1, \dots, R_{i-1} except, possibly, at n_i , and contains no edge that lies in one of R_1, \dots, R_{i-1} . Then, the set $\bigcup_{i \in \mathbb{N}} E(R_i)$ is thin, and its sum equals N since $\sum_{e \in E(R_i)} e = \{n_i\}$.

Suppose that $\mathcal{B} \subseteq \mathcal{N}$ is a basis of $\langle \mathcal{N} \rangle$, and let H be the graph with vertex set M and edge set \mathcal{B} . Since M is uncountable, so is \mathcal{B} . Therefore, one of the vertices in the countable set A , say v , is incident with infinitely many edges in \mathcal{B} . Denote by \mathcal{C} the set of those components in the graph obtained from H by deleting v (and its incident edges) that are adjacent to v .

Observe that for each $C \in \mathcal{C}$ there is exactly one edge in H between v and some vertex in C ; denote this vertex by u_C . Indeed, if there were two edges between v and C , we would easily find a cycle in \mathcal{B} , contradicting that \emptyset has a unique representation in \mathcal{B} since the sum of the edges of a cycle equals \emptyset .

Next, suppose there are distinct $C, D \in \mathcal{C}$ each containing a ray; then, C (respectively D) also contains a ray R (resp. S) starting at u_C (resp. at u_D).

Then $R \cup S$ together with the two edges between v and $\{u_C, u_D\}$ yields a set of edges which sums to \emptyset , again a contradiction.

Pick a countably infinite number of $C \in \mathcal{C}$ none of which contains a ray, and denote their set by \mathcal{C}' . As $N := \{u_C : C \in \mathcal{C}'\}$ is countable it lies in $\langle \mathcal{N} \rangle$, thus there is a $\mathcal{B}_N \subseteq \mathcal{B}$ such that $\sum_{e \in \mathcal{B}_N} e = N$.

Suppose there is a $C \in \mathcal{C}'$ such that an edge $e \in \mathcal{B}_N$ incident with u_C lies in C . As C does not contain any cycle or any ray, we can run from e along edges in $E(C) \cap \mathcal{B}_N$ to a vertex $w \neq u_C$ that is only incident with one edge in \mathcal{B}_N . This implies that $w \in \sum_{e \in \mathcal{B}_N} e = N$, a contradiction since $w \notin N$. However, u_C must be incident with an edge from \mathcal{B}_N . Consequently, for each $C \in \mathcal{C}'$ the edge between v and u_C lies in \mathcal{B}_N , contradicting that \mathcal{B}_N is thin. \square

5.3 Closedness

Let M be a set and R a ring. In general, for a subset \mathcal{N} of R^M , $\langle \mathcal{N} \rangle$ does not have to be closed under taking thin sums, even if R is a finite field. For instance, put $M = \mathbb{N}$, $R = \mathbb{Z}_2$ and consider $\mathcal{N} := \{\Phi(\{0, i\}) : i \in \mathbb{N}\}$. Easily, $\Phi(\mathbb{N}) \in \langle \langle \mathcal{N} \rangle \rangle$, but $\Phi(\mathbb{N}) \notin \langle \mathcal{N} \rangle$ as a representation in \mathcal{N} can only use finitely many of its elements. Thus, $\langle \mathcal{N} \rangle$ is not closed under taking thin sums. The critical property of this example is that \mathcal{N} contains infinitely many elements that “meet” the element 0 of M . Let us call a set $\mathcal{N} \subseteq R^M$ *thin* if for every $m \in M$, there are only finitely many elements N of \mathcal{N} such that $N(m) \neq 0$. We will show that $\langle \mathcal{N} \rangle$ is closed under taking thin sums if \mathcal{N} is thin and R is a field (Theorem 5.3) or a finite ring (Theorem 5.4). We will use the following easy lemma.

Lemma 5.1. *Let M be a set, R be a ring, and let \mathcal{T} be a thin set of elements of R^M . If $K \in \langle \langle \mathcal{T} \rangle \rangle$ and M' is a finite subset of M , then there is an element N of $\langle \mathcal{T} \rangle$ such that $N(m) = K(m)$ for every $m \in M'$.*

Proof. Let $d : \langle \mathcal{T} \rangle \rightarrow F$ be a representation of K in $\langle \mathcal{T} \rangle$, which exist as $K \in \langle \langle \mathcal{T} \rangle \rangle$. As M' is finite and d is thin, the set

$$T' = \{S \in \langle \mathcal{T} \rangle : d(S)S(m) \neq 0 \text{ for some } m \in M'\}$$

is finite.

Every element $S \in T'$ has a representation $a_S : \mathcal{T} \rightarrow F$ in \mathcal{T} . Define $a : \mathcal{T} \rightarrow F$, by

$$a(T) = \sum_{S \in T'} d(S)a_S(T). \quad (5.11)$$

We claim that $N := \sum_{T \in \mathcal{T}} a(T)T$ has the desired property. To show this, consider an $m \in M'$; we have to show that $\sum_{T \in \mathcal{T}} a(T)T(m) = K(m)$. Since d is a representation of K and since $d(S)S(m) = 0$ for $S \notin T'$, we obtain

$$K(m) = \sum_{S \in \langle \mathcal{T} \rangle} d(S)S(m) = \sum_{S \in T'} d(S)S(m).$$

By the definition of a_S , we have for every $S \in T'$ that $S(m) = \sum_{T \in \mathcal{T}} a_S(T)T(m)$. Replacing $S(m)$ in the above equation we get

$$\begin{aligned} K(m) &= \sum_{S \in T'} d(S) \sum_{T \in \mathcal{T}} a_S(T)T(m) = \sum_{S \in T'} \sum_{T \in \mathcal{T}} d(S)a_S(T)T(m) \\ &= \sum_{T \in \mathcal{T}} \sum_{S \in T'} d(S)a_S(T)T(m) = \sum_{T \in \mathcal{T}} a(T)T(m) = N(m), \end{aligned}$$

as desired, where in the last equation we used (5.11). \square

We can now proceed with the main results of this section.

Theorem 5.3 ([5]). *Let M be a set, F be a field, and let \mathcal{T} be a thin subset of F^M . Then $\langle \mathcal{T} \rangle$ is closed under taking thin sums, i.e. $\langle \mathcal{T} \rangle = \langle\langle \mathcal{T} \rangle\rangle$.*

Proof. Consider a $K \in \langle\langle \mathcal{T} \rangle\rangle$. We will reduce the problem of finding a representation of K in \mathcal{T} to the solution of an infinite system of equations. To do this, we associate a variable x^T with every $T \in \mathcal{T}$, and for each $m \in M$ we define e_m to be the linear equation

$$\sum_{T \in \mathcal{T}: T(m) \neq 0} x^T T(m) = K(m)$$

in the variables x^T . Let $E = \{e_m : m \in M\}$. By construction, if there is an assignment $a : \mathcal{T} \rightarrow F$ such that letting $x^T = a(T)$ for every $T \in \mathcal{T}$ yields a solution to every equation in E , then a is a representation of K in \mathcal{T} . So in what follows, our task is to find such a solution.

Let us start by claiming that every finite subset E' of E has a solution. Indeed, let M' be the set of $m \in M$ for which $e_m \in E'$. By Lemma 5.1, there is an $N \in \langle \mathcal{T} \rangle$ such that $N(m) = K(m)$ for every $m \in M'$, so let a' be a representation of N in \mathcal{T} . Then, letting $x^T = a'(T)$ for every $T \in \mathcal{T}$ yields a solution of E' .

For every $T \in \mathcal{T}$, define d_T to be the linear equation $x^T = 1$, and put $D = \{d_T : T \in \mathcal{T}\}$. Let E^* be a maximal subset of $E \cup D$ such that $E^* \supseteq E$ and every finite subset of E^* has a solution; to prove the existence of such an E^* , note that if $(E_i)_{i \in I}$ is a chain of nested supersets of E so that every

E_i has the property that every finite subset of E_i has a solution, then easily, their union $\bigcup_{i \in I} E_i$ also has this property. Thus by Zorn's lemma, there is a maximal set E^* with that property.

Next, we show that that for every $T \in \mathcal{T}$ there is a finite $E_T \subseteq E^*$ and an $f_T \in F$ such that $x^T = f_T$ in every solution of E_T . Suppose not. Then clearly, $d_T \notin E^*$. Consider a finite subset E' of E^* , and note that E' has at least one solution by the definition of E^* . If x^T takes the same value in all solutions of E' , then our assumption is contradicted as we can choose $E_T = E'$. Thus x^T takes two distinct values $x_1, x_2 \in F$ in two solutions of E' . Easily, for every $f \in F$ there is a solution of E' where $x^T = f' := fx_1 + (1 - f)x_2$. Letting $f = (1 - x_2)(x_1 - x_2)^{-1}$ we have $f' = 1$, thus there is a solution of E' in which $x^T = 1$, which means that $E' \cup \{d_T\}$ has a solution. Therefore, as E' was an arbitrary finite subset of E^* , every finite subset of $E^* \cup \{d_T\}$ has a solution, contradicting the maximality of E^* . Thus, E_T and f_T exist, as we have claimed.

Finally, define $a : \mathcal{T} \rightarrow F$ by $a(T) := f_T$. To see that a is a solution of E , consider an arbitrary $m \in M$. As \mathcal{T} is thin, $\mathcal{T}_m := \{T \in \mathcal{T} : T(m) \neq 0\}$ is finite. Thus, $E' := \{e_m\} \cup \bigcup_{T \in \mathcal{T}_m} E_T$ has, as a finite subset of E^* , a solution $b : \mathcal{T} \rightarrow F$. Since for every $T \in \mathcal{T}_m$, we have $E_T \subseteq E'$ it follows that $b(T) = f_T = a(T)$. As b solves e_m we see that a solves e_m , too. Thus a is a solution of E , and hence a representation of K in \mathcal{T} . \square

Theorem 5.4 ([5]). *Let M be a set, R be a finite ring, and let \mathcal{T} be a thin subset of R^M . Then $\langle \mathcal{T} \rangle = \langle\langle \mathcal{T} \rangle\rangle$.*

Proof. Let S be the set of ring elements of R , and consider the product space $X := \prod_{T \in \mathcal{T}} S = S^{\mathcal{T}}$ of $|\mathcal{T}|$ copies of the finite set S endowed with the discrete topology. By Tychonoff's theorem, X is a compact space. Its basic open sets have the form $O_b := \{a \in X : a|_U = b\}$, where b is some map from a finite set $U \subseteq \mathcal{T}$ to S .

Let $K \in \langle\langle \mathcal{T} \rangle\rangle$; we will show that $K \in \langle \mathcal{T} \rangle$. For every $m \in M$, let A_m be the set of elements $a \in X$ whose sum agrees with K on m ; that is, $\sum_{T \in \mathcal{T}} a(T)T(m) = K(m)$. As S is finite and there are only finitely many $T \in \mathcal{T}$ with $T(m) \neq 0$, it is easy to see that A_m is closed in X . By Lemma 5.1, for every finite $M' \subseteq M$, we have $\bigcap_{m \in M'} A_m \neq \emptyset$. By the finite intersection property of the sets A_m , their overall intersection is non-empty, and since, clearly, every element of $\bigcap_{m \in M} A_m \neq \emptyset$ is a representation of K in \mathcal{T} , we obtain $K \in \langle \mathcal{T} \rangle$. \square

So we have proved that $\langle \mathcal{T} \rangle$ is closed under taking this sums if \mathcal{T} is thin and R is a field or a finite ring. These results are best possible, since if R

is an infinite ring, then by the following example $\langle \mathcal{T} \rangle$ does not have to be closed.

Theorem 5.5 ([5]). *There is a thin $\mathcal{T} \subseteq \mathbb{Z}^{\mathbb{N}}$ such that $\langle \mathcal{T} \rangle \neq \langle\langle \mathcal{T} \rangle\rangle$.*

Proof. Define $N \in \mathbb{Z}^{\mathbb{N}}$ by $N(i) = 1$ for every $i \in \mathbb{N}$. For $j = 1, 2, \dots$, define $N_j \in \mathbb{Z}^{\mathbb{N}}$ by $N_j(j) = p_j$ and $N_j(i) = 0$ for every $i \neq j$, where p_j is the j -th prime number. Let $\mathcal{T} = \{N, N_1, N_2, \dots\}$, and note that \mathcal{T} is thin. We will show that the function $K \in \mathbb{Z}^{\mathbb{N}}$ defined by $K(i) = i$ is in $\langle\langle \mathcal{T} \rangle\rangle$ but not in $\langle \mathcal{T} \rangle$.

Let us first show that $K \notin \langle \mathcal{T} \rangle$. Suppose for contradiction, there is a representation $a : \mathcal{T} \rightarrow \mathbb{Z}$ of K in $\langle \mathcal{T} \rangle$. If $n := a(N) = 0$, then since $K = \sum_{L \in \mathcal{T}} a(L)L$ and $N_j(1) = 0$ for $j > 1$, we have $K(1) = \sum_{L \in \mathcal{T}} a(L)L(1) = a(N_1)N_1(1)$, which cannot be the case as $K(1) = 1$ and $N_1(1) = 2$. If $n > 0$, then $n + 1 = K(n + 1) = \sum_{L \in \mathcal{T}} a(L)L(n + 1) = a(N)N(n + 1) + a(N_{n+1})N_{n+1}(n + 1) = n + a(N_{n+1})p_{n+1}$, which implies that $1 = a(N_{n+1})p_{n+1}$ which cannot be the case as $p_{n+1} > 1$. Thus $n < 0$, but then we have for $n' = -n$ that $n' = K(n') = n + a(N_{n'})p_{n'}$, which means that $2n' \equiv 0 \pmod{p_{n'}}$, and since $p_{n'} > n'$ we have $p_{n'} = 2$ and $n = -1$. This again leads to a contradiction, as it implies that $K(3) = 3 = -1 + 5a(N_3)$.

Thus we showed that $K \notin \langle \mathcal{T} \rangle$ and we need to show that $K \in \langle\langle \mathcal{T} \rangle\rangle$. It suffices to show that for every $i \in \mathbb{N}$, there is an $S_i \in \langle \mathcal{T} \rangle$ so that $S_i(i) = 1$ and $S_i(j) = 0$ for every $j < i$ (we ignore $S_i(j)$ for $j > i$). Indeed, if such S_i exist, then they can represent K : we can inductively construct a representation d of K in $\{S_i : i \in \mathbb{N}\}$, starting by setting $d(S_1) = 1$, and letting $d(S_i) = i - \sum_{j < i} d(S_j)S_j(i)$.

So let us prove the existence of S_i . It suffices to find coefficients $a(N)$, $a(N_1)$, $a(N_2), \dots, a(N_i)$, such that

$$\begin{aligned} a(N) + a(N_i)p_i &= i, \text{ and} & (5.12) \\ a(N) + a(N_j)p_j &= 0 \text{ for every } 1 \leq j < i. \end{aligned}$$

Indeed, if we have such coefficients, then we can let $S_i = a(N)N + \sum_{j \leq i} a(N_j)N_j$. Now consider the system of congruences

$$\begin{aligned} x &\equiv i \pmod{p_i} \\ x &\equiv 0 \pmod{p_1} \\ x &\equiv 0 \pmod{p_2} \\ &\vdots \\ x &\equiv 0 \pmod{p_{i-1}} \end{aligned}$$

By the chinese remainder theorem, this system has a solution $n \in \mathbb{Z}$ for x . Now we can let $a(N) = n$, and solve each equation in (5.12) for the coefficient $a(N_j)$ that appears in it. We thus obtain the desired coefficients $a(N)$, $a(N_1)$, $a(N_2), \dots, a(N_i)$. This completes the proof. □

Chapter 6

Geodesic circles

6.1 Introduction

A finite cycle C in a graph G is called *geodesic* if, for any two vertices $x, y \in C$, the length of at least one of the two x - y paths on C equals the distance of x and y in G . It is easy to prove that (see Section 6.2.1):

Proposition 6.1. *The cycle space of a finite graph is generated by the circuits of its geodesic cycles.*

Although $\mathcal{C}(G)$ is known to be generated by the finite cycles of G , we shall see that the finite geodesic cycles need not generate $\mathcal{C}(G)$ —at least not as long as we measure path lengths the way we do for finite graphs, by counting edges. Indeed, when G is infinite then giving every edge length 1 will result in arc lengths that distort rather than reflect the natural geometry of $|G|$: edges ‘closer to’ ends should be shorter, if only to give arcs between ends finite lengths.

It looks, then, as though the question of whether or not Proposition 6.1 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 edge lengths are assigned, as long as the assignment satisfies a very general minimum requirement, the geodesic circles in $|G|$ — which we will promptly define — will generate $\mathcal{C}(G)$.

To make all this more precise, let us assume that when G was built as a 1-complex, the 1-cells used were real intervals of arbitrary lengths (instead of copies of the unit interval). Every arc in $|G|$ is, clearly, 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 (finite or infinite) sum of the lengths of these edges and half-edges. Given two points $x, y \in |G|$, write $d_G(x, y)$, or simply

$d(x, y)$ if G is fixed, for the infimum of the lengths of all x - y -arcs in $|G|$. If d_G is a metric on $|G|$ that induces its topology, call this 1-complex a *metric representation* of G . (In Section 6.3 we shall see that every locally finite connected graph has a metric representation.) We then call a circle C in $|G|$ *geodesic* if, for every two points $x, y \in C$, one of the two x - y -arcs in C has length $d_G(x, y)$. The set of edges of a geodesic circle in $|G|$ is a *geodesic circuit* in G . We can now state the main result of this chapter:

Theorem 6.1 ([23]). *For every metric representation of a connected locally finite graph G , $\mathcal{C}(G)$ is generated by the geodesic circuits in G .*

We prove Theorem 6.1 in Section 6.3, after showing that Proposition 6.1 holds finite graphs but not for infinite ones in Section 6.2. Finally, in Section 6.4 we will discuss some further problems.

In order to keep our expressions simple in this chapter, we will, with a slight abuse, not distinguish circles, paths and arcs from their edge sets.

6.2 Generating $\mathcal{C}(G)$ by geodesic cycles

6.2.1 Finite graphs

In this chapter finite graphs, like infinite ones, are considered as 1-complexes where the 1-cells (i.e. the edges) are real intervals of arbitrary lengths. We can thus define the *length* $L(X)$ of a path or cycle X in a finite graph G by $L(X) = \sum_{e \in E(X)} L(e)$, where $L(e)$ is the length of the edge e . A cycle C in G is, then, *L -geodesic*, if for any $x, y \in V(C)$ there is no x - y -path in G of length less than that of each x - y -path on C .

The following theorem generalises Proposition 6.1.

Theorem 6.2 ([23]). *For every finite graph G , every cycle C of G can be written as a sum of L -geodesic cycles of length at most $L(C)$.*

Proof. Suppose that the assertion is false for a graph G , and let D be a cycle in G of minimal length under all cycles C that cannot be written as a sum of L -geodesic cycles of length at most $L(C)$. As D is not L -geodesic, it is easy to see that there is a D -path P 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 -geodesic cycles of length less than $L(D)$, from which follows that D itself is such a sum, a contradiction. \square

By letting all edges have length 1, Theorem 6.2 implies Proposition 6.1.

6.2.2 Failure in infinite graphs

As already mentioned, Proposition 6.1 does not generalise to locally finite graphs. A counterexample is given in Figure 6.1. 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 $n \geq 2$, the n -th rung into $2n$ edges, we obtain H . The graph H is even a counterexample to the weaker assertion, that $\mathcal{C}(G)$ is generated by its *quasi-geodesic circles*, that is, circles that contain a shortest path between any two of their vertices (and thus contain at most one end). To see this, note that 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 must contain at least one rung, every quasi-geodesic circle (or geodesic cycle) contains e . On the other hand, Figure 6.2 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 quasi-geodesic circles to generate C , but since they all have to contain e the family cannot be thin.

The graph H is however no counterexample to Theorem 6.1 since if every edge has length 1, then d_H is not a metric of H .




Figure 6.1: A graph whose cycle space is not generated by its geodesic cycles.




Figure 6.2: An element of the topological cycle space (drawn thick) which is not the sum of a thin family of geodesic cycles.

6.3 Generating $\mathcal{C}(G)$ by geodesic circles

Metric representations do exist for every locally finite graph G . Just pick a normal spanning tree T of G with root $x \in V(G)$ (its existence is proved in [12, Theorem 8.2.4]), and define the length of any edge $uv \in E(G)$ as follows. If $uv \in E(T)$ and $v \in xTu$, let $L(uv) = \frac{1}{2^{|xTu|}}$. If $uv \notin E(T)$, let $L(uv) = \sum_{e \in uTv} L(e)$. It is easy to check that d is then indeed a metric of $|G|$ inducing its topology (see [10] for a proof).

In this section, let G be an arbitrary connected, locally finite graph, and consider a fixed metric representation of G . Every edge e has thus a length $L(e)$, and these edge lengths induce a metric $d = d_G$ of $|G|$ as defined in the Introduction.

The proof of Theorem 6.1 is not easy. It does not suffice to prove that every circle (or just every cycle) is a sum of a thin family of geodesic circles —in fact, the proof of the latter statement turns out to be as hard as the proof of Theorem 6.1. Although every element C of $\mathcal{C}(G)$ is a sum of a thin family of circles (even of cycles, see [12, Theorem 8.5.9]), a representation of every circle in this family as a sum of a thin family of geodesic circles will not necessarily yield such a representation for C , as the union of infinitely many thin families does not have to be thin.

In order to prove Theorem 6.1, we will use a sequence \hat{S}_i of finite auxiliary graphs whose limit is G . Given a $C \in \mathcal{C}(G)$ that we want to represent as a sum of geodesic circles, 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 geodesic cycles of \hat{S}_i , provided by Theorem 6.2. We will then use the resulting sequence of representations and compactness to obtain a representation of C as a sum of geodesic circles.

6.3.1 Restricting paths and circles

To define the auxiliary graphs mentioned above, pick a vertex $w \in G$, and for $i \in \mathbb{N}$ let $S_i = N^i(w)$; also let $S_{-1} = \emptyset$. Clearly, 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 with vertex set S_{i+1} , containing all edges of G incident with a vertex in S_i . Let \hat{S}_i be the graph resulting from \tilde{S}_i after joining each two vertices in $S_{i+1} - S_i$ that lie in the same component of $G - S_i$ with an edge, called an *outer edge*. For every $i \in \mathbb{N}$, L induces a length $L_i(e)$ to each edge e of \hat{S}_i : for every edge e that also lies in \tilde{S}_i , let $L_i(e) = L(e)$, and for every outer edge $e = uv$ of \hat{S}_i , let $L_i(e) = d(u, v)$.

For any arc or circle X in $|G|$ (respectively, for any path or cycle X in some $\hat{S}_j, j > i$), define the restriction $X|\hat{S}_i$ of X to \hat{S}_i to be the union of

$E(X) \cap \hat{S}_i$ with the set of all outer edges uv of \hat{S}_i such that X contains a u - v -arc (resp. u - v -path) having precisely its endpoints in common with \hat{S}_i , unless X contains no edge of \hat{S}_i , in which case let $X|\hat{S}_i = \emptyset$. 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.

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 6.1. *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 $L_i(X|\hat{S}_i) \leq L(X)$ (resp. $L_i(X|\hat{S}_i) \leq L_j(X)$).*

A consequence of this is the following:

Lemma 6.2. *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 $L_i(P) = d(x, y)$.*

Proof. Suppose first, that $L_i(P) < d(x, y)$. Replacing every outer edge uv in P by a u - v -arc of length $L_i(uv) + \varepsilon$ in $|G|$ for a sufficiently small ε , we obtain a topological x - y -path in $|G|$ shorter than $d(x, y)$. Since, by Lemma 2.5, the image of every topological path contains an arc with the same endpoints, this contradicts the definition of $d(x, y)$. Next, suppose that $L_i(P) > d(x, y)$. In this case, there is by the definition of $d(x, y)$ an x - y -arc Q in $|G|$ with $L(Q) < L_i(P)$, and since by Lemma 6.1 $L_i(Q|\hat{S}_i) \leq L(Q)$, $Q|\hat{S}_i$ is shorter than P contrary to our assumption. This completes the proof. \square

For the proof of Theorem 6.1, given a $C \in \mathcal{C}(G)$, we will construct a family of geodesic circles in ω steps, in each of which steps we will choose finitely many geodesic circles. To ensure that the resulting family will be thin, we will restrict the length of these circles. The following two lemmas will help us to do so. For every $i \in \mathbb{N}$, let

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

Note that as $|G|$ is compact, ε_i is finite.

Lemma 6.3. *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 -geodesic cycles in \hat{S}_j , each of which cycles has length at most $5\varepsilon_i$ with respect to L_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 , such that every cycle in this sum is either shorter than $5\varepsilon_i$ or a C -sector. Indeed, let Q be a C -part of D and let x, y be its endvertices. Every edge in Q has length at most $2\varepsilon_i$, because if e is an edge with length greater than $2\varepsilon_i$, then the middle-point 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 $3\varepsilon_i$. The distance of z from y along Q is also larger than ε_i . By the definition of ε_i , there is a z - y -path P in \hat{S}_j with $L_j(P) \leq \varepsilon_i$.

Let $Q_1 = zQy$ and let Q_2 be the other z - y -path in D . Clearly, $L(Q_2 + P) \leq 5\varepsilon_i$ (where $+$ denotes the symmetric difference), so $Q_2 + P$ can be written as a sum of edge-disjoint cycles in \hat{S}_j shorter than $5\varepsilon_i$. It is easy to see that $Q_1 + P$ can be written as a sum of C -sectors that are contained in $Q_1 + P$. As $Q_1 + P$ is shorter than D , each of those C -sectors is also shorter than D .

So every C -sector longer than $5\varepsilon_i$ is a sum of shorter cycles, either C -sectors or cycles shorter than $5\varepsilon_i$. Thus, as C is a C -sector itself, it is a sum of cycles not longer than $5\varepsilon_i$. By Lemma 6.2, every cycle in this sum is a sum of L_j -geodesic cycles in \hat{S}_j not longer than $5\varepsilon_i$; this completes the proof. \square

Lemma 6.4. *Let $\varepsilon > 0$ be given. There is an $n \in \mathbb{N}$, such that $\varepsilon_i < \varepsilon$ holds for every $i \geq n$.*

Proof. Suppose there is no such n . Thus, for every $i \in \mathbb{N}$, there is a component C_i of $G - S_i$ in which there are two points of distance at least ε . For every $i \in \mathbb{N}$, pick a vertex $c_i \in C_i$. As $|G|$ is compact (Lemma 2.3), and no vertex can be an accumulation point in $|G|$, there is an end ω in the closure of the set $\{c_0, c_1, \dots\}$ in $|G|$.

As $U := \{x \in |G| \mid d_L(x, \omega) < \frac{1}{2}\varepsilon\}$ is open in $|G|$, it has to contain $C(S_i, \omega)$ for some i . Furthermore, there is a vertex $c_j \in C(S_i, \omega)$ with $j \geq i$, because ω lies in the closure of $\{c_0, c_1, \dots\}$. As $S_j \supset S_i$, the component C_j of $G - S_j$ is contained in $C(S_i, \omega)$ 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 1. *Let $\varepsilon > 0$ be given. There is an $n \in \mathbb{N}$ such that for every $i \geq n$, every outer edge of \hat{S}_i is shorter than ε .*

6.3.2 Limits of paths and cycles

In this section we develop some tools that will help us obtain geodesic circles as limits of sequences of geodesic 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 $i \geq j$ is the restriction of X_{i+1} to \hat{S}_i .

Definition 6.1. *The limit of a chain X_j, X_{j+1}, \dots of paths or cycles, is the closure in $|G|$ of the set*

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

Unfortunately, the limit of a chain of cycles does not have to be a circle, as shown in Figure 6.3. However, we are able to prove the following lemma.




Figure 6.3: A chain X_0, X_1, \dots of cycles (drawn thick), whose limit X is no circle (but the disjoint union of two circles).

Lemma 6.5. *The limit of a chain of cycles is a continuous image of S^1 . The limit of a chain of paths is a continuous image of the real unit interval $[0, 1]$.*

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 $\sigma : S^1 \rightarrow X$ with the help of homeomorphisms $\sigma_i : S^1 \rightarrow X_i$ for every $i \in \mathbb{N}$. Start with some homeomorphism $\sigma_0 : S^1 \rightarrow X_0$. Now let $i \geq 1$ and suppose that $\sigma_{i-1} : S^1 \rightarrow X_{i-1}$ 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 $i \geq n$, 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 $\sigma : S^1 \rightarrow X$ 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 $\sigma(S^1) = X$.

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 $\sigma_i : [0, 1] \rightarrow X$ 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 $\sigma : [0, 1] \rightarrow X$ imitates that of the previous case. \square

Recall that a circle is geodesic, 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 geodesic if it contains no *shortcut*, that is, an arc in $|G|$ with endpoints $x, y \in C \cap \hat{V}$ and length less than both x - y -arcs 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.

Lemma 6.6. *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 no 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 4.1, \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 . \square

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

Lemma 6.7. *Let $i > j$ and let C be a geodesic circle in $|G|$ (respectively, an L_i -geodesic cycle in \hat{S}_i). Then $C_j := C|_{\hat{S}_j}$ is an L_j -geodesic 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 $L(Q_k) > d(x, y)$ (resp. $L_i(Q_k) > d(x, y)$) 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 $L_j(Q_k|_{\hat{S}_j}) > L_j(P)$. Moreover, by Lemma 6.1 we have $L(Q_k) \geq L_j(Q_k|_{\hat{S}_j})$ (resp. $L_i(Q_k) \geq L_j(Q_k|_{\hat{S}_j})$), and by Lemma 6.2 $L_j(P) \geq d(x, y)$, so our claim is proved. But then, by the definition of $d(x, y)$ (resp. by Lemma 6.2), there is an x - y -arc Q in $|G|$ such that $L(Q) < L(Q_k)$ (resp. an x - y -path Q in \hat{S}_i such that $d(x, y) = L_i(Q) < L_i(Q_k)$) for $k = 1, 2$, contradicting the fact that C is geodesic (resp. L_i -geodesic). \square

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

Lemma 6.8. *If C is the limit of a chain C_0, C_1, \dots of cycles, such that C_i is L_i -geodesic in \hat{S}_i , then C is a geodesic circle.*

Proof. Define the map σ as in the proof of Lemma 6.5 (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 $\sigma(P_2)$. By Corollary 1, 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 $i \geq j$, 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 -geodesic.

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 1, 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 -geodesic. Thus there is no shortcut to C , and therefore it is geodesic. \square

6.3.3 Proof of the generating theorem

Before we are able to prove Theorem 6.1, we need one last lemma.

Lemma 6.9. *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 a geodesic 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 $j \geq i$, choose, among all families \mathcal{H} of L_j -geodesic 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 $5\varepsilon_i$ with respect to L_j , and
- $\sum \mathcal{H}$ coincides with C in \tilde{S}_i .

By Lemma 6.3, 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.

Furthermore, for every $j \geq i$ and every $\mathcal{C} \in V_{j+1}$, restricting every cycle in \mathcal{C} to \hat{S}_j yields, by Lemma 6.7, a family \mathcal{C}^- of L_j -geodesic cycles. Moreover, \mathcal{C}^- lies in V_j : by Lemma 6.1, 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 $\bigcup_{j \geq i} V_j$, 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 $\mathcal{C}_i, \mathcal{C}_{i+1}, \dots$ such that for every $j \geq i$, $\mathcal{C}_j \in V_j$ and $\mathcal{C}_j = \mathcal{C}_{j+1}^-$. Therefore, for every cycle $D \in \mathcal{C}_i$, there is a chain $(D =) D_i, D_{i+1}, \dots$ of cycles such that $D_j \in \mathcal{C}_j$ for every $j \geq i$. By Lemma 6.8, the limit X_D of this chain is a geodesic 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 6.1. 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 $i = 0, 1, \dots$, we define finite families Γ_i of geodesic circles that satisfy the following condition:

$$C_i := C + \sum_{j \leq i} \bigcup \Gamma_j \text{ does not contain edges of } \tilde{S}_i \quad (6.1)$$

where $+$ denotes the symmetric difference.

By Lemma 2.1, there is a family \mathcal{C} of edge-disjoint circles whose sum equals C . Applying Lemma 6.9 to every circle in \mathcal{C} that meets S_0 (there are only finitely many), yields a finite family Γ_0 of geodesic circles that satisfies condition (6.1).

Now recursively, for $i = 0, 1, \dots$, suppose that $\Gamma_0, \dots, \Gamma_i$ are already defined finite families of geodesic circles satisfying condition (6.1), and write C_i as a sum of a family \mathcal{C} of edge-disjoint circles, supplied by Lemma 2.1. 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 6.9 yields a finite family \mathcal{F}_D of geodesic 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

$$C_{i+1} = C + \sum_{j \leq i+1} \bigcup \Gamma_j = C + \sum_{j \leq i} \bigcup \Gamma_j + \sum \Gamma_{i+1} = C_i + \sum \Gamma_{i+1}.$$

By the definition of Γ_{i+1} , condition (6.1) 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 $j \geq i$ no circle in Γ_j contains e . Indeed, there is an $i \in \mathbb{N}$, such that ε_j is smaller than $\frac{1}{5}L(e)$ for every $j \geq i$. Thus, by the definition of the families Γ_j , for every $j \geq i$, every circle in Γ_j is shorter than $L(e)$, 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 (6.1) and the claim above, there is an i , such that e is contained neither in C_i nor in a circle in $\bigcup_{j > i} \Gamma_j$. Thus, we have

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

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

6.4 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 [12, Corollary 8.5.9]). In the light of this result and Theorem 6.1, it is natural to pose the following question:

Problem 6.1 ([23]). *Let G be a locally finite graph, and consider a metric representation of G . Do the finite geodesic circles generate $\mathcal{C}(G)$?*

The answer to Problem 6.1 is negative: Figure 6.4 shows a graph with a metric representation where no geodesic circle is finite.

In the Introduction we did not define $d_G(x, y)$ 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 6.2 ([23]). *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 $\mathcal{P} = P_0, P_1, \dots$ be a sequence of x - y -arcs in $|G|$ such that $(L(P_j))_{j \in \mathbb{N}}$ converges to $d(x, y)$. Choose a $j \in \mathbb{N}$ such that every arc in \mathcal{P} meets S_j . Such a j always exists; if for example $x, y \in \Omega$, then pick j so that S_j separates every ray in x from every 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}, \dots is a chain of paths. The limit X of this chain contains x and y as it is closed, and $L(X) \leq d(x, y)$; for if $L(X) > d(x, y)$, then there is an i such that $L(X_i \cap \tilde{S}_i) > d(x, y)$, and as $L(X_k \cap \tilde{S}_k) > L(X_i \cap \tilde{S}_i)$ for $k > i$, this contradicts the fact that $(L(P_j))_{j \in \mathbb{N}}$ converges to $d(x, y)$. By Lemma 6.5, 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 $L(X) \leq d(x, y)$, 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 6.2 ([23]). *Let G be a locally finite graph, and consider a metric representation of G . 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 $L(P) - d(x, y) < \epsilon$?*

Surprisingly, the answer to this problem is also negative. The graph of Figure 6.4 with the indicated metric representation is again a counterexample.




Figure 6.4: A 1-ended graph G with a metric representation. Every geodesic circle is easily seen to contain infinitely many edges. Moreover, every (graph-theoretical) x - y -path has length at least 4, although $d(x, y) = 2$.

Theorems 6.2 and 6.1 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 assignment of edge lengths, as indicated by our next Problem. Call a cycle in a finite graph *peripheral*, if it is induced and non-separating.

Problem 6.3. *If G is a 3-connected finite graph, is there an assignment of lengths L to the edges of G , such that every L -geodesic cycle is peripheral?*

I was not able to give an answer to this problem. A positive answer would imply, by Theorem 6.2, a classic theorem of Tutte [31], asserting that the peripheral cycles of a 3-connected finite graph generate its cycle space. Problem 6.3 can also be posed for infinite graphs, using the infinite counterparts of the concepts involved¹.

¹Tutte's theorem has already been extended to locally finite graphs by Bruhn [2]

Chapter 7

Hamilton circles in squares of locally finite graphs

7.1 Introduction

We begin this chapter with a short proof of Fleischner's Theorem (Theorem 1.1) in a slightly stronger form (Section 7.2.2). Modifying this proof, we obtain a proof of a similar result of Fleischner [18], stating that the total graph of a connected, bridgeless finite graph contains a Hamilton cycle (Section 7.2.3). Then, we generalise Theorem 1.1 to locally finite graphs (Section 7.3). Such a generalisation has already been made by Thomassen, but only for locally finite graphs with one end:

Theorem 7.1 (Thomassen [29]). *If G is a 2-connected locally finite 1-ended graph, then G^2 contains a spanning double ray.*

Note that in a 1-ended graph, a double ray together with the end of the graph is a circle, thus Thomassen proved the existence of a Hamilton circle. Settling a conjecture of Diestel [11, 12], we will extend Fleischner's Theorem (and Thomassen's) to locally finite graphs with any number of ends:

Theorem 7.2 ([21]). *If G is a locally finite 2-connected graph, then G^2 has a Hamilton circle.*

As an intermediate step, we obtain a result which may be of independent interest. An *Euler tour* of G is a continuous map $\sigma : S^1 \rightarrow |G|$ that traverses every edge of G exactly once. An Euler tour is *end-faithful* if it visits each end of G exactly once. We will prove that:

Theorem 7.3 ([21]). *If a locally finite multigraph has an Euler tour, then it also has an end-faithful Euler tour.*

As discussed in Section 7.5, Theorem 7.3 could help generalise other sufficient conditions for the existence of a Hamilton circle.

7.2 Finite graphs

7.2.1 Definitions

A *trail* in a multigraph is a walk that does not traverse any edge more than once. A *pass* of a trail J through some vertex x , is a subtrail of J of the form $uexfv$ (where e and f are edges). If $P = uexfv$ is a pass of the trail J , then *lifting* P in J is the operation of replacing P by the trail ugv where g is a u - v edge if $u \neq v$, or by the trivial trail u if $u = v$ (in fact, the latter case will not occur), to obtain a new trail.

A *double edge* is a pair of parallel edges, and a *multipath* is a multigraph obtained from a path by replacing some of its edges by double edges. If $C \subseteq G$ are multigraphs, then a C -*trail* in G is either a path having precisely its endvertices (but no edge) in common with C , or a cycle having precisely one vertex in common with C . A vertex y on some cycle C is called C -*bound* if all neighbours of y lie on C .

7.2.2 Proof of Fleischner's theorem

We will use the following Lemma of Říha [32, 12]. For the convenience of the reader I'm including its proof.

Lemma 7.1. *If G is a 2-connected finite graph and $x \in V(G)$, then there is a cycle $C \subseteq G$ that contains x as well as a C -bound vertex $y \neq x$.*

Proof. As G is 2-connected, it contains a cycle C' that contains x . If C' is a Hamilton cycle there is nothing more to show, so let D be a component of $G - C'$. Assume that C' and D are chosen so that $|D|$ is minimal. Easily, C' contains a path P' between two distinct neighbours u, v of D whose interior \dot{P}' does not contain x and has no neighbour in D . Replacing P' in C' by a u - v -path through D , we obtain a cycle C that contains x and a vertex $y \in D$. By the minimality of $|D|$ and the choice of P' , y has no neighbour in $G - C$, so C satisfies the assertion of the lemma. \square

We will prove Theorem 1.1 in the following stronger form, which is similar to an assertion proved by Říha [32]. This proof is published in [22].

Theorem 7.4. *If G is a 2-connected finite graph and $x \in V(G)$, then G^2 has a Hamilton cycle both x -edges of which lie in $E(G)$.*

Proof. We perform induction on $|G|$. For $|G| = 3$ the assertion is trivial. For $|G| > 3$, let C be a cycle as provided by Lemma 7.1. Our first aim is to define, for every component D of $G - C$, a set of C -trails in $G^2 + E'$, where E' will be a set of additional edges parallel to edges of G . Every vertex of D will lie in exactly one such trail, and every edge of an element of such a trail that is incident with a vertex of C will lie in $E(G)$ or in E' .

If D consists of a single vertex u , we pick any C -trail in G containing u , and let E_D be the set of its two edges. If $|D| > 1$, let \tilde{D} be the (2-connected) graph obtained from G by contracting $G - D$ to a vertex \tilde{x} . Applying the induction hypothesis to \tilde{D} , we obtain a Hamilton cycle \tilde{H} of \tilde{D}^2 whose edges at \tilde{x} lie in $E(\tilde{D})$. Write \tilde{E} for the set of those edges of \tilde{H} that are not edges of G^2 . Replacing these by edges of G or new edges $e' \in E'$, we shall turn $E(\tilde{H})$ into the edge set of a union of C -trails. Consider an edge $uv \in \tilde{E}$, with $u \in D$. Then either $v = \tilde{x}$, or u, v have distance at most 2 in \tilde{D} but not in G , and are hence neighbours of \tilde{x} in \tilde{D} . In either case, G contains a u - C edge. Let E_D be obtained from $E(\tilde{H}) \setminus \tilde{E}$ by adding at every vertex $u \in D$ as many u - C edges as u has incident edges in \tilde{E} ; if u has two incident edges in \tilde{E} but sends only one edge e to C , we add both e and a new edge e' parallel to e . Then every vertex of D has the same degree (two) in $(V(G), E_D)$ as in \tilde{H} , so E_D is the edge set of a union of C -trails. Let $G' := (V(G), E(C) \cup \bigcup_D E_D)$ be the union of C and all those trails, for all components D together.

Let y be a C -bound vertex of C and pick a vertex z and edges f_1, f_2, g_1, g_2 of C , so that $C = xg_1z \dots f_1yf_2 \dots g_2x$. We will add parallel edges to some edges of $C - g_1$, to turn G' into an eulerian multigraph G_\emptyset — i.e. a connected multigraph in which every vertex has even degree (and which therefore has an Euler tour [12]). Every vertex in $G' - C$ already has degree 2. In order to obtain even degrees at the vertices in C , consider the vertices of C in reverse order, starting with x and ending with z . Let u be the vertex currently considered, and let v be the vertex to be considered next. Add a new edge parallel to uv if and only if u has odd degree in the multigraph obtained from G' so far. When finally $u = z$ is considered, every other vertex has even degree, so by the “hand-shaking lemma” z must have even degree too and no edge parallel to g_1 will be added. Let $C_\emptyset = G_\emptyset[V(C)]$.

If g_2 now has a parallel edge g'_2 , then delete both g_2, g'_2 . If g_2 has no parallel edge in G_\emptyset , and f_2 has a parallel edge f'_2 , then delete both f_2 and f'_2 . Let G_\emptyset be the resulting (eulerian) multigraph. If g_2 has been deleted, then let P_3 be the multipath $C_\emptyset - \{g_2, g'_2\}$. If not, let P_1 be the maximal multipath in C_\emptyset with endvertices x, y containing g_1 , and let P_2 be the multipath containing all edges in $E(C_\emptyset \cap G_\emptyset) - E(P_1)$.

Then, for every i such that P_i has been defined, do the following. Write $P_i = x_0^i x_1^i \dots x_{l_i}^i$ with $x_0^i = x$, and e_j^i or just e_j for the $x_{j-1}^i - x_j^i$ edge of P_i



Figure 7.1: Replacing e_j and e'_{j+1} by a new edge f_j .

in $E(C)$. Its parallel edge, if it exists, will again be denoted by e'_j (when i is fixed). Now for $j = 1, \dots, l_i - 1$, if e'_{j+1} is defined, replace e_j and e'_{j+1} by a new edge f_j joining x_{j-1} to x_{j+1} ; we say that f_j represents the trail $x_{j-1}e_jx_je'_{j+1}x_{j+1}$ (see Figure 7.1). Note that every such replacement leaves the current multigraph connected, and it preserves the parity of all degrees. Hence, the multigraph G^\triangleleft finally obtained by all these replacements is eulerian, so pick an Euler tour J of G^\triangleleft . Transform J into an Euler tour J' of G_\emptyset by replacing every edge in $E(J) - E(G_\emptyset)$ by the trail it represents.

Our plan is to perform some lifts in J' to transform it into a Hamilton cycle. To this end, we will now mark some passes for later lifting. Start by marking all passes of J' through x except for one arbitrarily chosen pass. We want to mark some more passes, so that for any vertex $v \in V(C) - x$ the following assertion is true:

$$\begin{aligned} &\text{if } v = x_j^i \text{ then all passes of } J' \text{ through } v \text{ are marked, except for} \\ &\text{the pass containing } e_j^i. \end{aligned} \quad (7.1)$$

This is easy to satisfy for $v \neq y$, as there is only one pair i, j so that $v = x_j^i$ in that case. A difficulty can only arise if $v = y = x_{l_1}^1 = x_{l_2}^2$, in case both P_1 and P_2 contain y . By the definition of the P_i this case only materialises if there are no edges g'_2, f'_2 in G_\emptyset , and as y is C -bound, it has degree at most 3 and hence degree 2 in G_\emptyset in that case. But then, there is only one pass of J' through v , which consists of $e_{l_1}^1, e_{l_2}^2$, and leaving it unmarked satisfies (7.1).

So we assume that (7.1) holds, and now we claim that

$$\begin{aligned} &\text{for every edge } e = uv \text{ in } J', \text{ at most one of the two passes of } J' \\ &\text{that contain } e \text{ was marked, and moreover if } u = x, \text{ then the pass} \\ &\text{of } J' \text{ through } v \text{ containing } e \text{ was not marked.} \end{aligned} \quad (7.2)$$

This is clear for edges in $E(G_\emptyset) - E(C_\emptyset)$, so pick an $e \in P_i$. If $e = e_j$ for some j , then by (7.1) the pass of J' through x_j^i containing e has not been marked; in particular, if e is incident with $x = x_0^i$, then $j = 1$ and the pass of J' through x_1^i containing e has not been marked. If $e = e'_j$, then e is not incident with x by the construction of G_\emptyset , and an edge f_{j-1} was defined to represent the trail $x_{j-2}e_{j-1}x_{j-1}e'_jx_j$. Since J contained f_{j-1} , this trail is a pass in J' . This pass is unmarked by (7.1), because it is a pass through x_{j-1} containing e_{j-1} .

So we proved our claim, which implies that no two marked passes share an edge. Thus we can now lift each marked pass of J' to an edge of G^2 , to

obtain a new closed trail H' in $G^2 + E'$ (in particular, H' is connected seen as a graph). Every vertex of G is traversed precisely once by H' , since by (7.1) we marked, and eventually lifted, for each vertex v of G all passes of J' through v except precisely one pass. (This is trivially true for a vertex u in $G - C$, as there is only one pass of J' through u and this pass was not marked.) In particular, H' cannot contain any pair of parallel edges, so we can replace every edge e' in H' that is parallel to an edge e of G by e to obtain a Hamilton cycle H of G^2 . Since by the second part of (7.2) no edge incident with x was lifted at its other end, both x -edges of H lie in G as desired. □

7.2.3 Total graphs

The *subdivision graph* G^\ddagger of a multigraph G is the bipartite graph with partition classes $V(G), E(G)$ where $x \in V(G)$ and $e \in E(G)$ are joined with an edge if x is incident with e in G . The *total graph* $T(G)$ of G is the graph on $V(G) \cup E(G)$ where two vertices are adjacent if the respective objects in G are adjacent or incident; equivalently, $T(G)$ is the square of G^\ddagger .

If $C \subseteq G$, then a C -cluster in G is either the union of a component D of $G - C$ with all C - D edges, or an edge that has both endvertices in C but does not lie in $E(C)$.

We will say that an edge uv of $T(G)$ *bridges* a vertex $x \in V(G)$ if u, v are x -edges of G .

By slightly modifying the proof in Section 7.2.2 we obtain a new proof of the following result of Fleischner [17, 18].

Theorem 7.5 ([17, 18]). *If G is a connected, bridgeless finite multigraph and $x \in V(G)$, then $T(G)$ has a Hamilton cycle H such that:*

- (i) *both x -edges of H lie in G^\ddagger and*
- (ii) *H contains no pair of incident edges e, f such that either both e, f bridge x or e bridges x and f is incident with x .*

Proof. We perform induction on $|G|$. For $|G| = 2$ the assertion is trivial. For $|G| > 2$, let C be a cycle in G containing x . Our first aim is to define, for every C -cluster D of $G - C$, a set of C -trails in $T(G)$. Every vertex of $T(D)$ not in C will lie in exactly one of these trails, and every edge of such a trail that is incident with a vertex of C will lie in $E(G^\ddagger)$.

If D is an edge e , then e^\ddagger is a C -trail with the desired properties. Let E_D be the edgeset of e^\ddagger . If D contains precisely one vertex u not in C , let

d_1, d_2, \dots, d_k be an enumeration of the u -edges of G ; since G is bridgeless, we have $k \geq 2$. Let p, q be the endpoints of d_1, d_k in C , and define the C -trail P in $T(G)$ by $P := pd_1d_2 \dots d_{k-1}ud_kq$. Note that the only edges of P incident with C are pd_1 and d_kq , and they lie in G^\ddagger as desired. Let E_D be the edgeset of P .

If D contains at least two vertices not in C , let \tilde{D} be the multigraph obtained from D by identifying all vertices in $C \cap D$ into a vertex \tilde{x} , and consider the following two cases.

If \tilde{D} is not bridgeless then, clearly, it contains a bridge b that is incident with \tilde{x} ; let u be the other vertex of b . Let d_1, d_2, \dots, d_k be an enumeration of the u - C edges in G ; since G is bridgeless, we have $k \geq 2$. Clearly, $\hat{D} := \tilde{D} - \tilde{x}$ is bridgeless, so applying the induction hypothesis to \hat{D} yields a Hamilton cycle $H =: uwH'w'u$ of $T(\hat{D})$ in which both u -edges uw, uw' lie in \hat{D}^\ddagger , thus also in G^\ddagger . We can now explicitly define a C -trail in $T(G)$ containing all vertices of $T(D)$ by $P := pd_1d_2 \dots d_{k-1}wH'w'ud_kq$, where p, q are the endvertices of d_1, d_k in C . Note that as $uw \in E(G^\ddagger)$, $d_{k-1}w$ is an edge of $T(G)$; moreover $pd_1, d_kp \in E(G^\ddagger)$. Let E_D be the edgeset of P .

Finally if \tilde{D} is bridgeless, then applying the induction hypothesis to \tilde{D} , we obtain a Hamilton cycle \tilde{H} of $T(\tilde{D})$ whose edges at \tilde{x} lie in \tilde{D}^\ddagger . Write \tilde{E} for the set of those edges of \tilde{H} that are not edges of $T(G)$. Replacing these by edges of G^\ddagger , we shall turn $E(\tilde{H})$ into the edge set of a union of C -trails. Consider an edge $uv \in \tilde{E}$, with $u \neq \tilde{x}$; then, either $v = \tilde{x}$ or u, v are \tilde{x} -edges in \tilde{D} . In the first case, $uv \in E(\tilde{D}^\ddagger)$ by (i), thus u is an edge of D incident with C , and there is a u - C edge $e(u)$ in G^\ddagger . In the second case, as u is an \tilde{x} -edge in \tilde{D} there is again a u - C edge $e(u)$ in G^\ddagger . By (ii), each vertex u of $T(\tilde{D}) - \tilde{x}$ is incident with at most one edge in \tilde{E} . Thus, if E_D is the edgeset obtained from $E(\tilde{H}) \setminus \tilde{E}$ by adding at every vertex $u \neq \tilde{x}$ that is incident with \tilde{E} the edge $e(u)$, then every vertex in $T(D) - C$ has the same degree (two) in $(V(T(G)), E_D)$ as in \tilde{H} , so E_D is the edge set of a union of C -trails, and any edge of such a C -trail incident with C lies in G^\ddagger .

Subdividing all edges of C and attaching to it all C -trails defined above, we obtain the (simple) graph

$$G' := (V(T(G)), E(C^\ddagger) \cup \bigcup_D E_D).$$

Let y, z be the neighbours of x in $C^\ddagger \subseteq G'$, and define the edges $g_1 := xy$ and $g_2 := xz$. We will add parallel edges to some edges of $C^\ddagger - g_1$, to turn G' into an eulerian multigraph G_δ . Every vertex in $G' - C^\ddagger$ already has degree 2. In order to obtain even degrees at the vertices in C^\ddagger , consider the vertices of C^\ddagger in its cyclic order, starting with x and ending with y . Let u be the vertex currently considered, and let v be the vertex to be considered

next. Add a new edge parallel to uv if and only if u has odd degree in the multigraph obtained from G' so far. When finally $u = y$ is considered, every other vertex has even degree, so by the “hand-shaking lemma” y must have even degree too and no edge parallel to g_1 will be added. Let G_{\emptyset} be the eulerian multigraph obtained so far, and let $C_{\emptyset} = G_{\emptyset}[V(C^{\pm})]$.

If g_2 has now a parallel edge g'_2 , then delete both g_2, g'_2 . Let G_{\emptyset} be the resulting (eulerian) multigraph. If g_2 has been deleted, then let P be the multipath $C_{\emptyset} - \{g_2, g'_2\}$. If not, let P be the multipath $C_{\emptyset} - g_1$. Write $P = x_0x_1 \dots x_{\ell_i}$ with $x_0 = x$, and e_j for the $x_{j-1}x_j$ edge of P in $E(C^{\pm})$. Its parallel edge, if it exists, will again be denoted by e'_j .

Our plan is to find an Euler tour J' of G_{\emptyset} that can be transformed into a Hamilton cycle of $T(G)$. In order to endow J' more easily with the required properties, we shall not define it directly. Instead, we shall derive J' from an Euler tour J of a related multigraph G^{\triangleleft} , which we define next.

For $j = 1, \dots, l_i - 1$, if e'_{j+1} exists, then replace e_j and e'_{j+1} by a new edge f_j joining x_{j-1} to x_{j+1} ; we say that f_j represents the trail $x_{j-1}e_jx_je'_{j+1}x_{j+1}$ (see Figure 7.1). Note that every such replacement leaves the current multigraph connected, and it preserves the parity of all degrees. Hence, the multigraph G^{\triangleleft} obtained from G_{\emptyset} by all these replacements is eulerian. So pick an Euler tour J of G^{\triangleleft} , and transform J into an Euler tour J' of G_{\emptyset} by replacing every edge in $E(J) - E(G_{\emptyset})$ by the trail it represents.

Our next aim is to perform some lifts in J' to transform it into a Hamilton cycle of $T(G)$. To this end, for every $j \geq 1$, mark all passes of J' through x_j that do not contain e_j . Also mark all passes of J' through x except for one pass (chosen arbitrarily).

We now claim that

for every edge $e = uv$ in J' , at most one of the two passes of J' that contain e was marked, and moreover if $u = x$, then the pass of J' through v containing e is unmarked. (7.3)

This is clear for edges in $E(G_{\emptyset}) - E(C_{\emptyset})$, so pick an $e \in E(C_{\emptyset})$. If $e \notin P$, then $e = xy$, and as $d_{G^{\pm}}(y) = 2$, there is only one pass of J' through y which, then, is unmarked, and contains e , so (7.3) holds for e . If $e \in P$, then $e = e_j$ or $e = e'_j$ for some j . If $e = e_j$, then the pass of J' through x_j containing e has not been marked; in particular, if e is incident with $x = x_0$, then $j = 1$ and the pass of J' through x_1 containing e has not been marked. If $e = e'_j$, then e is not incident with x , and an edge f_{j-1} was defined to represent the trail $x_{j-2}e_{j-1}x_{j-1}e'_jx_j$. Since J contained f_{j-1} , this trail is a pass in J' . This pass is unmarked, because it is a pass through x_{j-1} containing e_{j-1} . This completes the proof of (7.3).

Since all edges of G_{\emptyset} at vertices of C^{\dagger} lie in G^{\dagger} , all marked passes lift to edges of $T(G)$. As different marked passes never share an edge by (7.3), lifting them all at once turns J' into a closed trail H' in $T(G)$. Every vertex of $T(G)$ is traversed precisely once by H' , since we marked, and eventually lifted, for each vertex v of $T(G)$ all passes of J' through v except precisely one pass. (This is trivially true for a vertex u in $T(G) - C^{\dagger}$, as there is only one pass of J' through u and this pass was not marked.) In particular, H' cannot contain any pair of parallel edges, so we can replace every edge e' in H' that is parallel to an edge e of G^{\dagger} by e to obtain a Hamilton cycle H of $T(G)$. By the second part of (7.3), an edge of G_{\emptyset} incident with x can only be lifted at x , thus (i) holds. To see that (ii) holds too, note that if $e = uv$ is an edge in H that bridges x , then e resulted from lifting a pass $L = udx'd'v$ of J' . If in addition f is an edge in H incident with u such that f bridges x or $f = ux$, then this yields a pass of J' that shares no edges with L and contains an edge parallel to d . But x was not incident with any pair of parallel edges in G_{\emptyset} , a contradiction that proves (ii). □

7.3 Infinite graphs

7.3.1 Definitions

Let $G = (V, E)$ be a locally finite multigraph fixed throughout this section.

An x -edge is an edge incident with the vertex x .

A *shortcut* at a vertex x is the operation of replacing two edges ux, xv , where $u \neq v$, with a u - v -edge; the new edge *shortcuts* the edges ux, xv .

If $H \subseteq G$, then *contracting* H in G is the operation of replacing H in G with a new vertex z , and making z incident with all vertices of $G - H$ sending an edge to H . If G' is the graph resulting from G after contracting H to z , and $R \subseteq G'$, then $dc_z(R)$ is the subgraph of G resulting from R , after deleting z , in case $z \in V(R)$, and replacing each edge $xz \in E(R)$ with an arbitrarily chosen x - H -edge; you can think of $dc_z(R)$ as the result of decontracting z in R .

If $C \subseteq G$, denote by \hat{C} the union of C with all edges incident with C in G , including their endpoints. If $G \supseteq H$ and C is a component of $G - H$, then \hat{C} is called an H -bridge in G . Its *feet* are the vertices in $V(\hat{C}) - V(C)$.

A *multiedge* is the set of (parallel) edges between two fixed vertices of a multigraph. A *double edge* is a multiedge containing precisely two edges; a *single edge* is a multiedge containing precisely one edge. A *simple multigraph* is a multigraph all multiedges of which are either double or single edges.

If P is a path, $e \in E(P)$ and $x \in V(P)$, then xPe is the shortest subpath of P connecting x to an endvertex of e .

A *trail* in a multigraph is a walk in which no edge appears more than once.

An *Euler tour* of G is a continuous map $\sigma : S^1 \rightarrow |G|$ such that every inner point of an edge of G is the image of exactly one point of S^1 (thus, every edge is traversed exactly once, and in a “straight” manner). Call G *eulerian*, if it has an Euler tour. An *end-faithful* map $\sigma : S^1 \rightarrow |G|$ is a map such that every end in $\Omega(G)$ has exactly one preimage under σ .

The following lemma comes from [14, Theorem 7.2]. It has been proved for simple graphs only, but it is easily generalised to multigraphs.

Lemma 7.2. *The following three assertions are equivalent:*

- G is eulerian;
- $E(G) \in \mathcal{C}(G)$;
- Every finite cut of G is even.

7.3.2 Outline of the proof of Theorem 7.2

Before giving an outline of the proof of Theorem 7.2, let me compare it with the proof of Theorem 7.1 and the proof of Fleischner’s theorem by Říha [32, 12], which is shorter than its original proof. The descriptions that follow are approximate, omitting much information not needed for the comparison. Říha [32] proves Theorem 7.6 by induction; he finds a special cycle C , and then applies the induction hypothesis to every component of $G - C$, to obtain a set of C -paths in G^2 , called *basic paths*, so that each vertex of $G - C$ lies in exactly one of the paths. Basic paths have the property that their endedges are original edges of G ; let us call these endedges *bonds*. This property makes it possible, to recursively merge pairs of incident basic paths into longer basic paths, by shortcutting incident bonds, and he repeats this operation as often as possible without disconnecting the graph. Then, some edges of C are replaced by double edges, so that the resulting multigraph is eulerian. Finally, it is shown that every Euler tour J of this multigraph can be transformed to a Hamilton circle of G^2 , by replacing some subtrails of length two of J , with edges of G^2 with the same endvertices; we call this process the *hamiltonisation* of J .

Thomassen follows a similar plan in his proof of Theorem 7.1 (which appeared before Říha’s proof). The cycle C is replaced by a ray R , such that all components of $G - R$ are finite, and Theorem 7.6 is applied on each of

them, to give a set of R -paths in G^2 with the same properties as the basic paths in Ríha's proof. Then, some edges of R are duplicated, so that the resulting multigraph is eulerian. Next, some double edges are deleted, which splits R in finite paths, but does not disconnect the graph; let us call these paths *segments*. Again, some bonds are shortcutted, and it is then shown that every Euler tour J of this multigraph can be hamiltonised. Rather than doing the hamiltonisation on the whole graph simultaneously, it is shown that no matter how the restriction of an Euler tour J , to some segment and its neighbouring edges looks like, it is possible to locally modify J there, using edges of G^2 , so that it traverses each vertex of the segment exactly once. An example is shown in Figure 7.2.




Figure 7.2: An example of a local hamiltonisation. In the upper figure, the restriction of the Euler tour on the segment (horizontal path) is indicated; it consists of three paths. In the figure beneath, these trails have been transformed into disjoint paths in the square of the graph, that span all vertices (dashed lines).

Trying to imitate these proofs for arbitrary locally finite graphs, we face three major problems. The first one regards Euler tours. In the sketched proofs, an Euler tour was transformed to a Hamilton circle, by performing “leaps” over one vertex, using an edge of G^2 . Doing so for an arbitrary Euler tour of a locally finite graph, we cannot avoid running through some end more than once. But a Hamilton circle must, by definition, traverse each end exactly once, thus if we want to gain one from an Euler tour using this method, the Euler tour itself should be end-faithful. So we have to ask, which eulerian graphs admit an end-faithful Euler tour. The answer is given by Theorem 7.3: all of them.

The second problem is, what the analogue of R or C should be. In a graph with many ends there is no ray that leaves only finite components behind like the one used by Thomassen. Instead, we will use a complicated structure looking like a spanning tree of G containing two rays to each end, which spans the whole graph. Again, we will make the graph eulerian by duplicating edges, and we will split into finite segments. Like in Thomassen's proof, we want to make sure that we can change the chosen Euler tour locally on each segment W , so that it traverses each vertex exactly once. But in

order to be able to perform shortcuts with edges incident with W , as we did in Figure 7.2, we need an analogue of bonds: original edges of G , not affected by shortcuts performed while treating other segments. Indeed, we will make sure that the first edge of each segment W will not be shortcutted while treating W , so that other segments intersecting with W could shortcut it.

The third, and most serious problem, is that if we perform too many shortcuts we run the risk of changing the end topology of the graph. This problem appears even in the case of 1-ended graphs. Suppose, for example, that after performing the first steps of Thomassen's proof on some graph G having only one end ω , to find the ray R and the basic paths, we get the graph shown in Figure 7.3. If we shortcut every pair of incident bonds in this graph, we will end up with a 2-ended graph G' , because the basic paths will merge into a ray non-equivalent with R . We can still continue with the plan of finding an Euler tour and transforming it to a Hamilton circle H of G' , but H will not be a Hamilton circle of G : it will traverse ω twice.




Figure 7.3: Performing all shortcuts between bonds would create a new end.

Thomassen overcomes this difficulty by avoiding some shortcuts, at the cost of making the hamiltonisation of the Euler tour more difficult. See for example Figure 7.4, where vertex x is incident with two double edges on R , and two bonds. A possible restriction of an Euler tour on this segment is given, and the reader will confirm that it can only be locally hamiltonised in the way shown. Having two vertices like x on one segment can be fatal, as shown in Figure 7.5, where the Euler tour cannot be hamiltonised at all. But even one vertex like x on a segment is enough to cause problems; as already mentioned, we would like to hamiltonise each segment, so that its first edge is not shortcutted. This, however, is not possible in Figure 7.4. So on the one hand we should avoid shortcuts because they are dangerous for the end topology, and on the other we need them in order to get rid of vertices like x . An equilibrium is needed, which I could not find.

There is however an elegant solution to the problem, and it is achieved by imposing constraints on the Euler tour. These constraints specify trails of length 2 which the Euler tour must traverse. Technically, this is done by constructing an auxiliary graph, where each such trail has been replaced with an edge with the same endpoints. This auxiliary graph is eulerian if the original one is, and choosing an Euler tour of the auxiliary graph, and then




Figure 7.4: A difficult case: the lower figure shows the only possible hamiltonisation of the trails shown in the upper figure.




Figure 7.5: A case where no hamiltonisation is possible.

replacing the added edges with the trails they replaced, we obtain an Euler tour of the original graph, that indeed traverses the wanted trails. This is exemplified by Figure 7.6, which shows the auxiliary graph corresponding to the graph of Figure 7.5. Note that the problematic Euler tour of Figure 7.5 could not result from an Euler tour of the graph in Figure 7.6. It is this idea of imposing such constraints on the Euler tour that led to the short proof of Fleischner’s theorem presented in Section 7.2.2.




Figure 7.6: Applying constraints for an Euler tour of the graph of Figure 7.5. The dashed lines are edges of the original graph not in the auxiliary graph, while the continuous curved lines are edges of the auxiliary graph not in the original one.

The proof of Theorem 7.2 is structured as follows. We start by constructing the “scaffolding”, that is, the analogue of R in Thomassen’s proof, in Section 7.4.2. It consists of a set of ladder-like structures like the one shown in Figure 7.7 called *rope-ladders*, that are irregularly attached on each other, and a set of finite structures called *ear decompositions* that are attached on the rope-ladders. Unlike R , this scaffolding spans all vertices of our graph.

Next, we turn this scaffolding into an eulerian multigraph G^\emptyset by replacing some edges with double edges. Doing this is not as straightforward as in the finite case, and it will require its own section, Section 7.4.3. In Section 7.4.4, we will split G^\emptyset in segments, called larvae, as follows. We consider each Π

Thomassen are needed, with the exception of the articulation points. The reason we need shortcuts there is the following. It is no problem if two larvae W, W' intersect at the vertex where W' starts, since the hamiltonisation of the neighbourhood of W' will be done in such a way that its first edge is not affected, and then W will be able to shortcut this edge; if larvae intersect otherwise however, there could be a conflict between their hamiltonisations. As shown in Figure 7.8, it could happen that two larvae intersect at their last vertex, which is, in that case, an articulation point. In order to avoid a conflict, we make sure that if two larvae end at an articulation point y , then y has degree 2; if this is the case, then any Euler tour will traverse y only once, and therefore no conflict will arise during the hamiltonisation. Articulation points already existed in Říha's proof: there, C contains a vertex with the property that it sends no edge to the rest of the graph, and this vertex had a similar function. In infinite graphs however, it is not possible to pick articulation points without unwanted neighbours, but instead we will, in Section 7.4.5, perform shortcuts at the articulation points, to rid them of unwanted incident edges.

After doing all these changes, we are left with an auxiliary graph on $V(G)$, where we will, in Section 7.4.6, pick the end-faithful Euler tour. Then, based on the fact that the Euler tour complies with the constraints we imposed on it, and that the auxiliary graph bears the same end topology as the original one, we will show in Section 7.4.6 that it is possible to hamiltonise it to obtain a Hamilton circle of G^2 .

Summing up, the proof of Theorem 7.2 consists of the following steps:

- (i) construct the scaffolding;
- (ii) make it eulerian;
- (iii) split it into larvae;
- (iv) impose constraints on the Euler tour;
- (v) clean up the articulation points;
- (vi) pick an Euler tour and hamiltonise it.

7.3.3 Some preliminaries

End-devouring rays

The following lemmas are needed for the construction of the scaffolding. The graphs in Lemma 10 need not be locally finite, but the reader will lose

nothing by assuming that they are. Our definition of $\Omega(G)$ for arbitrary graphs remains that of Section 2.1.

If G is a graph and $\omega \in \Omega(G)$, we will say that a set K of ω -rays *devours* the end ω , if every ω -ray in G meets an element of K . An end devoured by some countable set of its rays will be called *countable*.

Lemma 7.3. *For every graph G and every countable end $\omega \in \Omega(G)$, if G has a set K of $k \in \mathbb{N}$ disjoint ω -rays, then it also has a set K' of k disjoint ω -rays that devours ω . Moreover, K' can be chosen so that its rays have the same starting vertices as the rays in K .*

Proof. We will perform induction on k . For $k = 1$ this is easy; the desired ray can for example be obtained by imitating the construction of normal spanning trees in [12, Theorem 8.2.4]. For the inductive step, let $K = \{R_0, R_1, \dots, R_{k-1}\}$ be a set of disjoint ω -rays in G . We want to apply the induction hypothesis to $G - R_0$, but we have to bear in mind that some of the R_i might not be equivalent to one another after deleting R_0 . So let $S \subset V$ be a finite set, so that any two tails of elements of K that lie in the same component of $G - R_0 - S$ are equivalent. Applying the induction hypothesis to all components of $G - R_0 - S$ that contain a tail of some R_i , we obtain a new set of rays $R'_1, R'_2, \dots, R'_{k-1}$ so that any ray equivalent to some R_i in $G - R_0 - S$ meets some R'_j , and for each j , R'_j starts at the first vertex of R_j not in S . We can now prolong each R'_j using the subpath of R_j that lies in S , to achieve that R'_j and R_j start at the same vertex (without loss of generality, each R_j leaves S only once, because otherwise we can add an initial subpath of R_j to S). Moreover, let R'_0 be a ray in $G - \{R'_1, R'_2, \dots, R'_{k-1}\}$ meeting all rays equivalent with R_0 in that graph and starting at the first vertex of R_0 . We claim that $K' = \{R'_0, R'_1, \dots, R'_{k-1}\}$ meets every ω -ray in G .

Indeed, suppose that $S \in \omega$, $S \cap \bigcup K' = \emptyset$, and let \mathcal{P} be a set of infinitely many disjoint S - R_0 -paths in G . Now either infinitely many of these paths avoid $\{R'_1, R'_2, \dots, R'_{k-1}\}$, or infinitely many meet the same R'_i before meeting R_0 . In the first case, S is equivalent with R_0 in $G - \bigcup \{R'_1, R'_2, \dots, R'_{k-1}\}$, and thus meets R'_0 , whereas in the second case, S is equivalent with some R'_i in $G - R_0 - S$ and thus meets some R'_j , a contradiction that proves the claim. \square

Lemma 7.4. *If G is locally finite, $\omega \in \Omega$, and K is a set of rays devouring ω in G , then every component of $G - K$ sends finitely many edges to K .*

Proof. If such a component sends infinitely many edges to K , then by Lemma 2.6 it contains a comb whose spine is equivalent with the rays in K , contradicting the fact that K meets every ω -ray. \square

End-faithful topological Euler tours

In this section we prove Theorem 7.3:

Proof of Theorem 7.3. By Lemma 7.2, every finite cut of G is even. Then G has a finite cycle C , because otherwise, every edge would form a cut. Let $\sigma_0 : S^1 \rightarrow C$ be a continuous function, that maps a closed interval of S^1 to each vertex and edge of C (think of the edges as containing their endvertices).

We will now inductively, in ω steps, define an end-faithful topological Euler tour σ of G . After each step i , we will have defined a finite set of edges F_i , and a continuous surjection $\sigma_i : S^1 \rightarrow \bar{F}_i$, where \bar{F}_i is the subspace of $|G|$ consisting of all edges in F_i and their incident vertices. In addition, we will have chosen a set of vertices S_i incident with F_i , and for each $v \in S_i$ a closed interval I_v of S^1 mapped to v by σ_i (These intervals will be used in subsequent steps to accommodate the rest of the graph). Then, at step $i+1$, we will pick a suitable set of finite cycles from $E(G) - F_i$, put them in F_i to obtain F_{i+1} , and modify σ_i to σ_{i+1} . We might also add some vertices to S_i to obtain S_{i+1} .

Formally, let $F_0 = E(C)$, $S_0 = \emptyset$ and σ_0 as defined above. Let e_1, e_2, \dots be an enumeration of the edges of G . Then, perform ω steps of the following type (skip 0). At step i , let for a moment $S_i = S_{i-1}$ and consider the components of $G - F_{i-1}$. For each of them, say D , there is, by construction, at most one vertex $v \in S_i$ incident with D . If there is none, just pick any vertex v incident with both D and F_{i-1} , put it in S_i and let I_v be any of the closed intervals of S^1 mapped to v by σ_{i-1} . Furthermore, pick a finite cycle C_D in D incident with v . Try to choose C_D so that it contains e_j , where $j = \min_{e_k \in E(D)} k$, and if it is not possible, choose C_D so that the distance between e_j and C_D is smaller than the distance between e_j and F_{i-1} . Then, to define σ_i , map I_v continuously to C_D , mapping an initial and a final closed subinterval of I_v to v , and a closed subinterval of I_v to each vertex and edge of C_D , and let all these subintervals have equal length. Redefine I_v to be one of those subintervals that were mapped to v .

We claim that the images $\sigma_i(x)$ of each point $x \in S^1$ converge to a point in $|G|$. Indeed, since $|G|$ is compact, it suffices to show that $(\sigma_i(x))_{i \in \mathbb{N}}$ cannot contain two subsequences converging to different points. It is easy to check that if $(\sigma_i(x))_{i \in \mathbb{N}}$ contains a subsequence converging to a vertex or an inner point of an edge, then $(\sigma_i(x))_{i \in \mathbb{N}}$ also converges to that point. So suppose it contains two subsequences converging to two ends ω, ω' , and find a finite edge set F separating those ends. Note that $F \subset F_i$ for i large enough, so denote by D, D' the components of $G - F_i$ that contain rays of ω, ω' respectively. But if x is mapped on a point p by σ_{i+1} , then for all steps succeeding $i+1$, x will

be mapped on a point belonging to the component of $G - F_i$ that contains p . Thus $(\sigma_i(x))_{i \in \mathbb{N}}$ cannot meet both D, D' for $n > i$, a contradiction that proves the claim.

So we may define

$$\begin{aligned} \sigma : S^1 &\rightarrow |G|, \\ x &\mapsto \lim_{n \rightarrow \infty} \sigma_n(x) \end{aligned}$$

In order to prove that σ is continuous, we have to show that the preimage of any basic open set of $|G|$ is open. This is obvious for basic open sets of vertices and inner points of edges. For every $\omega \in \Omega$, the sequence of basic open sets of ω that arise after deleting $F_i, i \in \mathbb{N}$, is converging, so it suffices to consider the basic open sets of that form, and it is easy to see that their preimages are indeed open.

Thus σ is continuous, and by the way we chose the C_D , it traverses each edge exactly once, which makes it an Euler tour.

We now claim that every end $\omega \in \Omega$ has at most one preimage under σ . Since at every step i , there is only one vertex v in S_i meeting the component of $G - F_i$ that contains rays of ω , I_v is the only interval of S^1 in which ω could be accommodated. Since I_v gets subdivided after every step, the claim is true, and thus σ is end-faithful. \square

7.4 Proof of Theorem 7.2

7.4.1 A stronger assertion

A shorter proof of Fleischner's theorem was given by Říha [32], who in fact proved a slightly stronger assertion:

Theorem 7.6. *Let G be a finite 2-connected graph, let $y^* \in V(G)$ and let $e^* = y^*x^* \in E(G)$. Then, G^2 has a Hamilton cycle H that contains e^* and a y^* -edge $e' \in E(G)$ with $e' \neq e^*$.*

Rather than Theorem 1.1, we will generalise this stronger assertion:

Theorem 7.7 ([21]). *Let G be an infinite 2-connected locally finite graph, let $y^* \in V(G)$ and let $e^* = y^*x^* \in E(G)$. Then, G^2 has a Hamilton circle H that contains e^* and a y^* -edge $e' \in E(G)$ with $e' \neq e^*$.*

7.4.2 Constructing the scaffolding

In this section we construct the “scaffolding” G^\sharp mentioned in Section 7.3.2. The scaffolding will be made of two ingredients: *rope-ladders* and *ear decompositions*. Let us give the definition of the latter and some motivation.

An *ear decomposition* of a finite H -bridge B in G , where $H \subseteq G$, is a subgraph of B spanning $V(B - H)$, that consists of a sequence C_1, C_2, \dots, C_n of paths called *ears*, C_i having the distinct endvertices p_i, q_i , so that

- C_1 is an H -path, i.e. $C_1 \cap H = \{p_1, q_1\}$;
- $C_i \cap (H \cup \bigcup_{j < i} C_j) = \{p_i, q_i\}$ for every i ;
- C_i is not an H -path for $i > 1$, and
- for every i , C_i contains a vertex $y(C_i) \neq p_i, q_i$ all of whose neighbours in G lie in $H \cup \bigcup_{j \leq i} C_j$ (thus $|C_i| \geq 2$).

The endedges of C_i are its *bonds*, and $y(C_i)$ is its *articulation point*. An ear decomposition is what we get from the special cycle C in the proof of Říha (see Section 7.3.2) if we try to make a constructive proof out of Říha’s inductive one. To see this, recall that in that proof, after choosing C we applied the induction hypothesis to every component D of $G - C$. To be more precise, the induction hypothesis is in fact not applied to D , but to an auxiliary graph \tilde{D} resulting from G after deleting all components of $G - C$ other than D and contracting C to a vertex z . If we wish to have a constructive proof, we can start the procedure again, with \tilde{D} instead of G : we can choose a special cycle $C' \ni z$ in \tilde{D} , as we chose C in G , and so on. Now if we decontract z , C' will look like an arc of an ear decomposition.

The special cycle C in Říha’s proof contained a special vertex, and articulation points play the role of that vertex.

The role of the ear decompositions in our proof will be to take care of finite pieces of G that are not in any rope-ladder. The following lemma is similar to a lemma of Říha [32].

Lemma 7.5. *If $G \supseteq H$ is a 2-connected graph, B is a finite H -bridge, and x is a foot of B , then B has an ear decomposition such that x lies in C_1 .*

Proof. Pick an H -path C starting at x , and let D be a component of $B - (C \cup H)$; if there is no such component, then we can let $C_1 = C$, pick any inner vertex of C_1 as $y(C_1)$, and choose C_1 as an ear decomposition of B . Suppose that C, D have been chosen so that $|V(D)|$ is minimal. Clearly, D has at least one neighbour u on $C - H$. If it has more than one, then let P be a subpath of $C - H$ whose endvertices u, v are neighbours of D , such that

no inner vertex of P is a neighbour of D , and let C_1 be the path resulting from C after replacing P with a v - u -path with all inner vertices in D . If u is the only neighbour of D on $C - H$, then let v be a neighbour of D in H , and replace one of the subpaths of C connecting u to H , with a v - u -path with all inner vertices in D (having at least one inner vertex), so that the resulting path C_1 meets x .

In both cases, C_1 contains a vertex $y \in D$, and we can let $y(C_1) = y$, because if y had a neighbour in $B - (C_1 \cup H)$, it would lie in a component $D' \subsetneq D$ of $B - (C_1 \cup H)$, contradicting the choice of C, D .

Now for $i = 2, 3, \dots$, suppose that C_1, C_2, \dots, C_{i-1} have already been defined and satisfy the conditions imposed by the definition of an ear decomposition on its ears. If there is a vertex u of $B - H$ they do not contain, let $H' = H \cup \bigcup_{j < i} C_j$, and repeat the above procedure for the H' -bridge B' that contains u , but this time letting a foot of B' in $H' - H$ play the role of x (this makes sure that C_i is not an H -path), to define the path C_i . If there is no such vertex u , then C_1, C_2, \dots, C_{i-1} is the wanted ear decomposition. \square

Let us now turn our attention to rope-ladders. Rope-ladders are in a way similar to ear decompositions; Their pis (see Figure 7.7) play a similar role as the ears of an ear decomposition, although there are important differences. For example, we cannot guarantee that all neighbours in G of any articulation point lie in its pi, but instead we will, as explained in Section 7.3.2, perform shortcuts so as to rid articulation points of unwanted neighbours. In order to be able to perform these shortcuts without changing the end topology, we have to pick articulation points far enough from each other. But this will be an easy task, because we can choose the pis to be arbitrarily long.

By a result of Halin [24, Theorem], if G is a locally finite 2-connected graph, then there are for any $v \in V(G)$ and any $\omega \in \Omega(G)$, two independent ω -rays starting at v . If $x, y \in V(G)$, then by applying this result on ω and an imaginary vertex joined to both x, y with an edge, we obtain the following:

Lemma 7.6. *In a locally finite 2-connected graph G , there are for any $x, y \in V(G)$ and any $\omega \in \Omega(G)$, two disjoint ω -rays starting at x, y respectively.*

Let ω be any end of G . By Lemma 7.6, there are two disjoint ω -rays starting at x^* and y^* (recall that x^*, y^* are the special vertices in the assertion of Theorem 7.7), and by Lemma 7.3 there are rays R^0, L starting at y^*, x^* respectively, that devour ω . Let $L^0 = y^*x^*L$, and let $r_0^0, l_0^0 = y^*$. Choose a sequence $(y_j^0)_{j \in \mathbb{N}}$ of vertices of R^0 , and a sequence $(P_j^0)_{j \in \mathbb{N}}$ of pairwise disjoint R^0 - L^0 paths, P_j^0 having the endpoints r_{j+1}^0, l_{j+1}^0 , so that y_0^0 is the first vertex on R^0 after r_1^0 , and for each $j > 0$ the following conditions are satisfied (see Figure 7.7):

- y_j^0 lies on $\mathring{y}_{j-1}^0 R^0$;
- r_{j+1}^0 lies in $\mathring{y}_j^0 R^0 \mathring{y}_{j+1}^0$, and l_{j+1}^0 lies in $\mathring{l}_j^0 L^0$;
- Every $(R^0 \cup L^0)$ -bridge that has y_{j-1}^0 as a foot, has all other feet on $\mathring{r}_{j-1}^0 R^0 \mathring{r}_{j+1}^0 \cup \mathring{l}_{j-1}^0 L^0 \mathring{l}_{j+1}^0 - y_0^0$

(The last condition makes sure that the articulation points are “far” from each other.) All these conditions are easy to satisfy, if we choose the y_j^0 and P_j^0 in the order $P_0^0, y_1^0, P_1^0, y_2^0, P_2^0, \dots$: recall that by Lemma 7.4 every $(R^0 \cup L^0)$ -bridge has only finitely many feet, so each time we want to choose a new y_j^0 or P_j^0 , we just have to go far enough along R^0 and L^0 .

Let RL^0 be the subgraph of G consisting of $RL_-^0 := R^0 \cup L^0 \cup \{P_j^0 | j \in \mathbb{N}\}$ and the ears of a fixed ear decomposition of every finite RL_-^0 -bridge, which exists by Lemma 7.5. Let $\mathring{RL}^0 = RL^0 - y^*$ and let $G_0^\# = RL^0$.

The construction of $G_0^\#$ was the first step in an infinite procedure the aim of which is to define $G^\#$. Each step i of this procedure will be similar to the construction of $G_0^\#$: we will choose rays R^i, L^i in $G - G_{i-1}^\#$, and add them together with some R^i - L^i -paths and some ear decompositions to $G_{i-1}^\#$ to obtain $G_i^\#$. The endpoints of R^i, L^i will be distinct vertices of $G_{i-1}^\#$.

Formally, let $(x_i)_{i \in \mathbb{N}}$ be an enumeration of $V := V(G)$, and perform ω steps of the following type, skipping step 0. At step i , let C_i be the component of $G - G_{i-1}^\#$ containing x_j , where j is the smallest index so that $x_j \notin G_{i-1}^\#$; if no such j exists, then stop the procedure and set $G^\# = G_{i-1}^\#$. If the path Q_j has not been defined yet, then let it be any x_j - \mathring{RL}^l -path in \hat{C}_i , where l is the greatest index for which such a path exists. Let $v = v(i)$ be the last vertex of Q_j not in $G_{i-1}^\#$, and $w = w(i)$ the vertex after v on Q_j (thus $w \in G_{i-1}^\#$).

Intuitively, we want to have x_j in $G_i^\#$, but this might be impossible if x_j is “far” from $G_{i-1}^\#$, in which case we just try to make sure that $G_i^\#$ is closer to x_j than $G_{i-1}^\#$ was. In order to make “closer” precise, we define the path Q_j , and in each subsequent step we eat up part of Q_j till we reach its endpoint x_j ; later we will formally prove that this does work. The condition that Q_j meet $G_{i-1}^\#$ at \mathring{RL}^l is needed in order to guarantee that $G^\#$ has the same end topology as G . To see why this condition should help retain the topology, it is useful to compare with the construction of a normal spanning tree. Recall that as seen in Chapter 3, a normal spanning tree of a locally finite graph G has the same end topology as G . A normal spanning tree can be constructed by starting with the root and no edges, and stepwise attaching new vertices to the already constructed tree, but each new vertex has to be attached as high as possible on the existing tree (see [12]). The construction

of G^\sharp imitates this, in the sense that rope-ladders are stepwise attached on each other, and the aforementioned condition on Q_j expresses the fact that new rope-ladders should be attached “as high as possible”.

We claim that:

Claim 7.1. *There are disjoint rays $R^i \approx L^i$ in \hat{C}_i , starting at w, G_{i-1}^\sharp respectively, that devour some end of G , so that either $v \in R^i \cup L^i$ or v lies in a finite component of $C_i - R^i \cup L^i$.*

Proof. Contracting $G - C_i$ to one vertex z , we obtain a 2-connected graph, in which we can apply Lemma 7.6 and Lemma 7.3 to get disjoint rays $R' \approx L'$, starting at v and z respectively, that devour some end of C_i (C_i is infinite because at the end of each step i we add all finite components to G_i^\sharp). By Lemma 7.4, C_i has finitely many feet, thus R', L' also devour some end of G . If $L^* := dc_z(L')$ does not start at w , then $R^i := wvR', L^i := L^*$ satisfy the conditions of the claim. If L^* does start at w , then let P be a G_{i-1}^\sharp - $(R' \cup L^*)$ -path in $G - w$. If the endpoint u of P lies on L^* (respectively R'), then let $R = wvR', L = PuL^*$ (respectively $R = PuR', L = L^*$). In the first case (if $u \in L^*$), $v \in R \cup L$ holds so we can choose $R^i = R, L^i = L$.

In the second case, we can suppose that R', L', P have been chosen so that the path $W := wvR'u$ is minimal. Now if v lies in R or in a finite component of $C_i - R \cup L$ we can again choose $R^i = R, L^i = L$. Otherwise, we may contract $G_{i-1}^\sharp \cup R \cup L$ to a vertex z' , and as above, find disjoint rays $R'' \approx L''$, starting at v and z' respectively, that devour some end of G . We distinguish two cases:

If $L^{**} := dc_{z'}(L'')$ meets W , let r (respectively l) be the last vertex of R'' (L^{**}) on W (note that $r \neq u$). Now if $r \in lWu$, let $R^i = RuWrR''$ and $L^i = wWlL^{**}$, whereas if $l \in rWu$, let $R^i = RuWlL^{**}$ and $L^i = wWrR''$. Depending on whether $l = w$ or not, R^i, L^i either contradict the minimality of W , or contain v and thus satisfy all conditions of the Claim.

If L^{**} does not meet W , then there are three subcases. In the first subcase, L^{**} starts at L . Then, let v' be the last vertex on W meeting R'' , and choose $L^i = LL^{**}, R^i = RuWv'R''$. In the second subcase, L^{**} starts at R , and we can choose $L^i = RL^{**}, R^i = wvR''$, and in the third subcase, L^{**} starts at G_{i-1}^\sharp , and we can choose $L^i = L^{**}, R^i = wvR''$. Depending on whether $v = v'$ or not, R^i, L^i either contain v and thus satisfy all conditions of the Claim, or contradict the minimality of W . □

With $R^i =: r_0^i R^i, L^i =: l_0^i L^i$ having been chosen as in the Claim, pick a sequence $(y_j^i)_{j \in \mathbb{N}}$ of vertices of R^i , and a sequence $(P_j^i)_{j \in \mathbb{N}}$ of pairwise disjoint R^i - L^i paths in C_i , P_j^i having the endpoints r_{j+1}^i, l_{j+1}^i , so that y_0^i is the first

vertex on R^i after the endpoint of P_0^i , and for each $j > 0$ the following conditions are satisfied:

- y_j^i lies on $\hat{y}_{j-1}^i R^i$;
- r_{j+1}^i lies in $\hat{y}_j^i R^i \hat{y}_{j+1}^i$, and l_{j+1}^i lies in $\hat{l}_j^i L^i$;
- Every $(G_{i-1}^\# \cup R^i \cup L^i)$ -bridge in G that has y_{j-1}^i as a foot, has all other feet on $\hat{r}_{j-1}^i R^i \hat{r}_{j+1}^i \cup \hat{l}_{j-1}^i L^i \hat{l}_{j+1}^i - y_0^i$

Such a choice is possible because, by Lemma 7.4 every $(G_{i-1}^\# \cup R^i \cup L^i)$ -bridge in G has finitely many feet, and there are only finitely many $(G_{i-1}^\# \cup R^i \cup L^i)$ -bridges in G with feet on both $G_{i-1}^\#$ and $R^i \cup L^i$ (again, we choose the y_j^i and P_j^i in the order $P_0^i, y_1^i, P_1^i, y_2^i, P_2^i, \dots$).

Let RL^i be the graph consisting of $RL_-^i := R^i \cup L^i \cup \{P_j^i | j \in \mathbb{N}\}$ and the ears of a fixed ear decomposition of every finite RL_-^i -bridge in G . We call RL^i a *rope-ladder* (RL^0 is also a rope-ladder). Let $\hat{R}L^i = RL^i - \{r_0^i, l_0^i\}$. Recall that one of R^i, L^i contains an edge incident with w . Call this edge the *anchor* of RL^i , unless $w = y_j^k$ for some j, k , in which case let the other edge of $R^i \cup L^i$ incident with $G_{i-1}^\#$ be the anchor of RL^i (by the choice of the articulation points, it cannot be the case that both these edges are incident with some articulation point). Note that by the choice of Q_j and of the y_j^i , the anchor of RL^i is incident with RL^l , where l is the highest index so that C_i has a foot on RL^l . We will say that RL^i is *anchored* on RL^l . Call the edge $e^* = y^*x^*$ the *anchor* of RL^0 .

Define the relation \prec between rope-ladders, so that $R \prec R'$ if R' is anchored on R , and let \preceq be the reflexive transitive closure of \prec . Clearly, \preceq is a partial order.

For every $i \geq 0, j \geq 1$, call the cycle in RL_-^i containing P_j^i, P_{j-1}^i a *window* of RL^i , and denote it by W_j^i . Moreover, let Π_0^i denote the path $r_0^i R^i P_0^i L^i l_0^i$, and for any $j \geq 1$, let $\Pi_j^i = W_j^i - \Pi_{j-1}^i$. For every $i, j \in \mathbb{N}$, call Π_j^i a *pi*, let $y(\Pi_j^i) = y_j^i$, and call y_j^i an *articulation point*. The *bonds* of a pi are its endedges. The *bonds* of W_j^i are the bonds of Π_j^i and the *bonds* of RL^i are the bonds of Π_0^i (that is, its endedges). Call the edges of RL_-^0 incident with y^* the *bonds* of RL^0 . Recall that ears also have bonds and articulation points. The following assertion is true by construction:

Observation 1. *If RL^i sends a bond to $\hat{R}L^j$, then $RL^j \preceq RL^i$.*

For suppose that RL^i sends a bond to $\hat{R}L^j$ but $RL^j \not\preceq RL^i$. Since RL^j must have been constructed before RL^i , we have $j < i$, and thus $RL^i \not\preceq RL^j$.

Let k be the greatest index such that $RL^k \preceq RL^i$ and $RL^k \preceq RL^j$ (this is well-defined as $RL^0 \preceq RL^i, RL^j$). Clearly, $RL^k \neq RL^i, RL^j$. Now if RL^i, RL^j lie in the same RL^k -bridge C in G , then by the choice of the paths Q_j , $R \preceq RL^i, RL^j$ holds, where R is the first rope-ladder constructed in C , and R contradicts the choice of RL^k . Thus RL^i, RL^j lie in distinct RL^k -bridges, contradicting the fact that RL^i sends a bond to $\mathring{R}L^j$.

Similarly, we can prove that:

Observation 2. *If an ear decomposition sends edges to $\mathring{R}L^i$ and $\mathring{R}L^j$ then either $RL^j \preceq RL^i$ or $RL^i \preceq RL^j$.*

For every i , let the *anchor* of Π_0^i be the anchor of RL^i . For every Π_j^i with $j > 0$, pick one of its bonds and call it its *anchor*. For any ear of an ear decomposition, pick one of its bonds that is not incident with any y_j^i and call it its *anchor*.

Define the relation \prec between pis and ears (we are using, with a slight abuse, the same symbol for two relations) so that $\Pi \prec \Pi'$ if either $\Pi = \Pi_j^i$ and $\Pi' = \Pi_{j+1}^i$ for some i, j , or $\Pi' = \Pi_0^i$ and RL^i sends a bond to an inner vertex of Π for some i , or Π' is an ear and it sends a bond to an inner vertex of Π (y^* is an inner vertex of Π_0^0). Let \preceq be the reflexive transitive closure of \prec . Clearly, \preceq is a partial order.

Define G_i^\sharp as the union of G_{i-1}^\sharp with RL^i and an ear decomposition of every finite $(G_{i-1}^\sharp \cup RL^i)$ -bridge.

We can now define $G^\sharp := \bigcup_{i \in \mathbb{N}} G_i^\sharp$. In the rest of the paper we will be working with this graph instead of G , but in order to be able to do so we have to show that it does not differ from G too much.

Let us prove that $V(G^\sharp) = V$. By the choice of R^i, L^i , either $v(i) \in R^i \cup L^i$ or $v(i)$ lies in a finite component of $C_i - R^i \cup L^i$. In both cases, $v(i) \in G_i^\sharp$. Thus, at most $|Q_j|$ steps after the path Q_j is defined, x_j will lie in G_i^\sharp , which implies that $V(G^\sharp) = V$.

Our next aim is to prove that $|G^\sharp| \cong |G|$, and we will do so using Lemma 2.9. Suppose there are rays Q, T in G^\sharp such that $Q \not\approx_{G^\sharp} T$ but $Q \approx_G T$. They could not belong to the end of R^i for any i , because then they would have to meet RL^i infinitely often, and thus, clearly, be equivalent in G^\sharp . Thus there is a j so that G_j^\sharp separates a tail of Q from a tail of T in G^\sharp (just choose j large enough so that G_j^\sharp contains some finite Q - T -separator). We will show that this is not possible. Indeed, since $Q \approx_G T$, there is a component C of $G - G_j^\sharp$ containing tails of both Q, T . Clearly, Q has some vertex in C that lies on some RL_-^i , and the same holds for T . So pick $k, l \in \mathbb{N}$ so that $q \in V(Q) \cap C \cap RL_-^k$ and $t \in V(T) \cap C \cap RL_-^l$. If R is the first rope-ladder constructed in C , then by the choice of the paths Q_i ,

$R \preceq L$ holds for any rope-ladder L meeting C , in particular $R \preceq RL^k, RL^l$. Thus, we can find a $t\text{-}\mathring{R}$ -path P_1 in G^\sharp that uses only vertices of rope-ladders RL^i such that $R \preceq RL^i \preceq RL^l$, and a $q\text{-}\mathring{R}$ -path P_2 in G^\sharp , that uses only vertices of rope-ladders RL^i such that $R \preceq RL^i \preceq RL^k$. But P_1, P_2 and \mathring{R} lie in C , contradicting the fact that G_j^\sharp separates Q from T in G^\sharp .

Thus no such rays Q, T exist and by Lemma 2.9, $|G| \cong |G^\sharp|$.

7.4.3 Making the graph eulerian

The next step is to replace some edges of G^\sharp with double edges, in order to turn it into an eulerian simple multigraph G_i^\emptyset , but so that no anchor is replaced with a double edge. Rather than constructing the simple multigraph explicitly, we will show its existence using Theorem 2.1. In order to meet its condition, we will show that:

Claim 7.2. *For every $i \in \mathbb{N}$ there is an eulerian simple multigraph G_i^\emptyset on V , so that any two vertices are neighbours in G_i^\emptyset if and only if they are neighbours in G^\sharp , and furthermore no anchor that lies in $G^\sharp[y^*]^i$ — that is, the subgraph of G^\sharp induced by the vertices of distance at most i from y^* — is replaced with a double edge in G_i^\emptyset .*

Proof. If C is a cycle of length at least 3 in the simple multigraph G , then *switching* C is the operation of replacing in G each single edge of C with a double edge, and each double edge containing an edge of C with a single edge. Note that switching a cycle in a simple multigraph does not affect vertex degrees.

In order to prove the Claim, begin by doubling all edges of G^\sharp . Then, for every ear decomposition C_1, C_2, \dots, C_k meeting $G^\sharp[y^*]^i$, recursively, for $j = k, k-1, \dots, 0$, if the anchor of C_j is now a double edge, find a cycle containing C_j and avoiding $\bigcup_{i>j} C_i$ and all other ear decompositions in G^\sharp , and switch this cycle. After doing so for all ear decompositions, recursively for $j = l, l-1, \dots, 1$, where l is the greatest index such that the anchor of RL^l lies in $G^\sharp[y^*]^i$, if the anchor of RL^j is a double edge, switch a cycle comprising Π_0^j and a path in G_{j-1}^\sharp that has the same endvertices as Π_0^j and contains no edge of an ear decomposition (for $j = 0$ switch Π_0^0). After the end of this recursion, switch every window whose anchor is a double edge and lies in $G^\sharp[y^*]^i$.

Let G_i^\emptyset be the resulting simple multigraph. Note that G_i^\emptyset resulted from a simple multigraph where all multiedges are double, after switching a finite set of cycles. Since switching a cycle does not affect the parity of a finite cut, G_i^\emptyset is eulerian by Lemma 7.2. Obviously, G_i^\emptyset satisfies all conditions of the Claim. \square

In order to apply Theorem 2.1, define for every edge $e \in G^\sharp$ a logical variable $v(e)$, the truth-values of which encode the two possible multiplicities of e , and let \mathcal{V} be the set of these variables. For every finite cut F of G^\sharp , write a propositional formula with variables in \mathcal{V} , expressing the fact that the sum of the multiplicities of the edges in F is even. Moreover, for every anchor e in G^\sharp , write a propositional formula with the only variable $v(e)$, expressing the fact that e is not replaced with a double edge.

By Theorem 2.1 and the Claim, there is an assignment of truth-values to the elements of \mathcal{V} satisfying all these propositional formulas. This assignment encodes an assignment of multiplicities to the edges of G^\sharp , which defines a simple multigraph G^\emptyset which is eulerian (by Lemma 7.2), and in which all anchors of G^\sharp form single edges.

Let G^\emptyset be the simple multigraph resulting from G^\emptyset after deleting each double edge that has the same endvertices as a bond in G^\sharp . Trivially, $|G^\emptyset| \cong |G^\sharp|$ holds, and we claim that furthermore $|G^\emptyset| \cong |G^\emptyset|$. In order to prove this assertion, we will specify a thin set of detours for the deleted edges and apply Lemma 2.10.

If $e = pq$ is a deleted bond of a rope-ladder RL^i , let $j = \min(\{k < i \mid e \cap V(RL^k) \neq \emptyset\})$, and suppose that p lies in RL^j and q in $\mathring{R}L^i$. We claim that there is a p - q -path $dt(e)$ in G^\emptyset that satisfies the following conditions:

- (i) all vertices of $dt(e)$ lie in rope-ladders R such that $RL^j \preceq R \preceq RL^i$;
- (ii) $dt(e)$ avoids all pis of RL^j below the first one that sends an edge to the component C of $G^\sharp - G_j^\sharp$ that contains $\mathring{R}L^i$.

To prove this, note that as each pi lost at most one bond and no other edges, $\mathring{R}L^l \cap G^\emptyset$ is connected for every l , and so if RL^l is anchored on RL^k , then for any vertex r in $\mathring{R}L^l$ there is an r - $\mathring{R}L^k$ -path in $RL^l \cap G^\emptyset$ that uses the anchor of RL^l . As $RL^j \preceq RL^i$ by Observation 1, we can use this fact recursively to obtain a q - RL^j -path P in $G^\emptyset \cap G_i^\sharp$ all of whose vertices lie in rope-ladders R such that $RL^j \preceq R \preceq RL^i$, and thus avoids the pis of RL^j below the first one that sends an edge to C . Since p lies in a pi that sends an edge to C (namely, pq), and each pi lost at most one bond, we can prolong P by a path in $RL^j \cap G^\emptyset$ to obtain the desired path $dt(e)$.

If e is a deleted bond of an ear of an ear decomposition D , let i be the greatest index such that $\mathring{R}L^i$ meets D and let j be the least index such that $\mathring{R}L^j$ meets D . Then, Observation 2 yields $RL^j \preceq RL^i$, and we can, by a similar argument as in the previous case, find a p - q -path $dt(e)$ in G^\emptyset all vertices of which lie in D and in rope-ladders R such that $RL^j \preceq R \preceq RL^i$ and which avoids the pis of RL^j below the first one that meets the

$(G_{j-1}^\# \cup RL_-^j)$ -bridge in $G^\#$ that contains D . Finally, if e is a deleted bond of a window W , then let $dt(e) = W - e$.

We claim that the set $\{dt(e) | e \in E(G^\emptyset) - E(G^\#)\}$ is thin. To prove this, it suffices to show that for any fixed edge f , there are only finitely many rope-ladders and ear decompositions that can contribute a $dt(e)$ containing f . This is clear if f lies in an ear decomposition D , as a detour $dt(e)$ can only go through f in that case if $e \in D$ (see (i)), so suppose that $f \in \Pi_m^l$ for some l, m . Let us start by showing that there are finitely many rope-ladders that contribute a $dt(e)$ containing f . By (i) there are two kinds of rope-ladders R that have a deleted bond e such that $dt(e)$ contains f : the ones for which e meets $\mathring{R}L^l$ (i.e. the rope-ladder containing f), and the ones for which e meets some rope-ladder $L \preceq RL^l$, $L \neq RL^l$. By (ii), the rope-ladders of the first kind belong to components of $G^\# - G_l^\#$ that send an edge to some Π_k^l with $k \leq m$. Since the graph is locally finite, there are only finitely many such components, and by Lemma 7.4 each of them sends finitely many edges to RL^l . As these edges are the only candidates for e , there are only finitely many rope-ladders of the first kind. Let R be a rope-ladder of the second kind, let e be its deleted bond, and let RL^k be the rope-ladder on which RL^l is anchored ($RL^l \neq RL^0$ by the definition of the second kind). Then by (i), R and $\mathring{R}L^l$ lie in the same component C of $G^\# - G_k^\#$, and as e has to be one of the edges between C and $G_k^\#$, which again by Lemma 7.4 are only finitely many, there are only finitely many rope-ladders of the second kind that can contribute a $dt(e)$ containing f .

It remains to show that there are finitely many ear decompositions that contribute a $dt(e)$ containing f . To see this, note that any such ear decomposition D must lie in a $(G_{l-1}^\# \cup RL_-^l)$ -bridge in $G^\#$ that has feet in both $G_{l-1}^\# \cup \bigcup_{k \leq m} \Pi_k^i$ and RL_-^l by the definition of $dt(e)$, and by the construction of $G^\#$ there are only finitely many such bridges; indeed, any such bridge lies in the $G_{l-1}^\#$ -bridge in G in which RL^l lies, and this bridge has finitely many feet on $G_{l-1}^\#$ by Lemma 7.4. Again, every $(G_{l-1}^\# \cup RL_-^l)$ -bridge in $G^\#$ sends finitely many edges to $(G_{l-1}^\# \cup RL_-^l)$ by Lemma 7.4, and as D must send an edge to $G_{l-1}^\# \cup \bigcup_{k \leq m} \Pi_k^i$ by the definition of $dt(e)$, there are finitely many ear decompositions that contribute a $dt(e)$ containing f .

This proves our claim that the set $\{dt(e) | e \in E(G^\emptyset) - E(G^\#)\}$ is thin, which by Lemma 2.10 implies that $|G^\emptyset| \approx |G^\emptyset| \approx |G^\#|$.

7.4.4 Splitting into finite multigraphs

While constructing $G^\#$ we defined many useful terms like pis, windows, rope-ladders, etc. that were subgraphs of $G^\#$. We want to use those names and

symbols for G^\emptyset as well, but these symbols now do not denote the original graphs, but the simple multigraphs in G^\emptyset that replaced them. Thus, when referring to G^\emptyset , we will use Π_j^i to denote the subgraph of G^\emptyset spanned by the multiedges whose endvertices were joined by an edge of Π_j^i in G^\sharp , a bond (respectively anchor) is a multiedge whose endvertices were joined by a bond (resp. anchor) in G^\sharp , and so on. Moreover, xy denotes the multiedge with endvertices x, y .

According to our plan, as stated in Section 7.3.2, we want to split the graph in larvae; let us introduce them formally. A *larva* is a pair (s, P) , where P is a *multipath* — i.e. a simple multigraph obtain from a path after replacing some of its edges with double edges — in G^\emptyset , s is one of its endvertices, called its *mouth*, and the multiedge of P incident with s is a single edge. For every larva $W = (s, P)$, we label the vertices of P with $x_i = x_i(W)$, so that $P = x_0(=s)x_1x_2 \dots x_n$. Moreover, let $e_i = e_i(W)$ denote the multiedge $x_{i-1}x_i$, and if e_i is a double edge denote its edges by e_i^-, e_i^\cap , otherwise let e_i^- be its only edge. Let $P(W) = P$. Whenever we use an expression assuming a direction on P or W , we consider x_0 to be its first vertex and x_n its last. In order to simplify the notation, we will also write sPy instead of (s, sPy) .

Recall that we want to impose some constraints on the Euler tour that is supposed to produce a Hamilton circle of G . This is done separately for each larva following the pattern of Figure 7.6: *metamorphosing* the larva $W = (s, P)$, is the operation of replacing, in P and in G^\emptyset , the edges e_{i+1}^\cap, e_i^- , for every i such that e_{i+1} is a double edge, by an $x_{i-1}x_{i+1}$ edge f_i , called a *representing* edge (representing edges already existed in our proof of Fleischer's theorem, and they had the same function there). Note that e_{i+1}^\cap, e_i^-, f_i form a triangle. The *caterpillar* of W is the graph X resulting from P after metamorphosing W . Note that X is connected. Each time we metamorphose a larva, we will assume that for each deleted edge e , a detour $dt(e)$ for e is specified in X .

If P has length at least 2 and the last multiedge e_k of P is a single edge, then *completely metamorphosing* W is the operation of metamorphosing W , and then replacing e_k^-, e_{k-1}^- with an $x_{k-2}x_k$ edge f_{k-1} , also called a *representing* edge. If W is completely metamorphosed, then its *pseudo-mouth* is its last vertex. The *double caterpillar* of W is the graph X resulting from P after completely metamorphosing W . A double caterpillar has a big advantage in comparison to a caterpillar: the additional constraint (on the Euler tour), allows it to be hamiltonised so that its last edge, as well as its first, is not shortcutted, and so its pseudo-mouth is allowed to meet other larvae (even if it is not an articulation point). This advantage however, comes at a high price: a double caterpillar is a disconnected graph, with two components. For this reason, each time we completely metamorphose a larva W to obtain

X , we will specify some *detour* $dt(X)$ for X , that is, an X -path connecting the two components of X (note that the last two vertices of $P(W)$ lie in distinct components of X , and in fact $dt(X)$ will always be a path connecting those vertices). We assume that also for each edge e deleted while completely metamorphosing W to get X , a detour $dt(e)$ for e is specified in $X \cup dt(X)$.

We now divide the graph in larvae, and either metamorphose or completely metamorphose each of them. (According to the sketch of the proof in Section 7.3.2, we first split the graph into larvae and then impose the constraints on the Euler tour, but in fact these two steps will be performed simultaneously, the constraints being imposed by metamorphosing or completely metamorphosing the larvae.) Formally, we will specify a set of edge-disjoint larvae \mathcal{W} so that $G^\emptyset = \bigcup_{W \in \mathcal{W}} P(W)$, and the following conditions are satisfied:

Condition 1. *If $W, W' \in \mathcal{W}$, then W, W' are edge-disjoint, and if v is a vertex lying in both W and W' then one of the following is the case:*

- *v is the mouth of W or W' ;*
- *v is the pseudo-mouth of W or W' ; or*
- *v is an articulation point, both W, W' end at v , and the last multiedges of both W, W' are single (none of W, W' will be completely metamorphosed in this case).*

Condition 2. *For every $v \in V - y^*$, there is an element $W(v)$ of \mathcal{W} containing v , so that v is neither the mouth nor the pseudo-mouth of $W(v)$ (by Condition 1, there is at most one $W \in \mathcal{W}$ with this property, unless v is an articulation point).*

We will construct a simple multigraph G^\triangleleft on V by performing operations of the following kinds on G^\emptyset :

- replacing two edges e, f with an edge forming a triangle with e, f ;
- switching a window;
- adding a double edge from $G^\emptyset - G^\emptyset$;
- deleting a double edge.

Note that metamorphosing or completely metamorphosing a larva is a set of operations of the first kind. Each time we delete an edge, we will specify a detour in G^\triangleleft , so as to be able to use Lemma 2.10 to prove that we did not

change the end topology. The fact that we only use the above operations will imply that the graph remains eulerian after all changes.

Define \mathcal{W} to be the set of larvae that we will metamorphose or completely metamorphose in what follows. For any pi or ear Π , denote by $a(\Pi)$ the endvertex of Π incident with its anchor, and by $b(\Pi)$ the other endvertex of Π ($a(\Pi_0^0) = b(\Pi_0^0) = y^*$).

In Section 7.3.2, and in particular in Figure 7.8, the rules according to which we split the graph in larvae were roughly given. The idea behind these rules, is to keep the graph induced by $V(RL_-^i)$ connected for every i , so as to guarantee that the end topology remains the same. If however, we apply those rules to Π_0^i , then we could disconnect part of it from the rest of RL_-^i . To avoid this, we will treat pis of the form Π_0^i differently.

So we will construct G^\triangleleft in two phases, in the first of which we will take care of the pis of the form Π_0^i , and in the second of the rest of the graph. At any point of the construction it will be an easy check — left to the reader — that Condition 1 holds for all larvae defined up to that point. Moreover, each pi or ear Π will be considered at some point, and then every vertex in $\Pi - \{a(\Pi), b(\Pi)\}$ will be put in some larva in \mathcal{W} without being its mouth or its pseudo-mouth. As $a(\Pi), b(\Pi)$ lie in some other pi or ear as well, this is enough to guarantee that Condition 2 will be satisfied.

For the first phase, perform ω steps of the following kind. In step i , if Π_0^i has already been handled, that is, divided in larvae, in some previous step, or if one of its bonds is not present in G^\emptyset , go to the next step. Otherwise, if zw was a bond of Π_1^i such that there is no $z-w$ -edge in G^\emptyset , then add a $z-w$ double edge. We consider two cases.

In the first case, called Case I, both multiedges $e = r_1^i y, e'$ incident with $y := y_0^i$ on Π_1^i are single or both are double edges. If they are both double, then switch W_1^i . No matter if we switched W_1^i or not, metamorphose the larvae (r_1^i, e) (this is a trivial larva) and $l_0^i \Pi_0^i l_1^i \Pi_1^i y$ (see Figure 7.9). If the multiedge $d = l_1^i l'$ of P_0^i incident with l_1^i is double, delete d and metamorphose the larva $r_0^i \Pi_0^i l'$; pick a detour $dt(d)$ for d in the three resulting caterpillars. If d is single, and there is a double edge f on P_0^i , delete f and metamorphose the larvae $l_1^i P_0^i f$ and $r_0^i \Pi_0^i f$; pick a detour $dt(f)$ in the four resulting caterpillars. If there is no double edge on P_0^i , let r' be the neighbour of r_1^i on P_0^i , metamorphose the larva $l_1^i P_0^i r'$ and completely metamorphose the larva $r_0^i \Pi_0^i r'$ (see Figure 7.10); a detour for the double caterpillar can be found in the resulting caterpillars. It is easy to confirm that the following is true:

Observation 3. *No detour specified in Case I meets any pi $\Pi \neq \Pi_0^i$ for which $\Pi \preceq \Pi_0^i$ holds.*

Observation 3 and other observations of this kind that will follow will

help prove that the set of detours that will be defined in this section is thin.




Figure 7.9: Splitting into larvae: Case I, and d is double. The dashed lines indicate larvae, and arrows show away from the mouth.




Figure 7.10: Splitting into larvae: Case I, and no double edge on P_0^i . The line with arrows at both ends indicates a larva that will be completely metamorphosed.

In the second case, called Case II, one of e, e' is single and the other is double. We want to choose and metamorphose some larvae, that will give rise to an $\mathring{R}L_-^i$ -path A^i in G^{\triangleleft} with one endpoint at y , which will help us delete an edge in RL_-^i without putting the end topology at risk; A^i will help by being part of a detour for the deleted edge.

Since y has even degree in G^{\emptyset} , there is at least one single bond (other than e) incident with y . Pick such a bond b , so that the pi or ear Π_0 of which b is a bond is minimal with respect to \preceq . Note that the anchor of Π_0 cannot be b , since y is an articulation point. Let Π_1 be the pi or ear that contains $a(\Pi_0)$ as an inner vertex. Metamorphose the larva $(a(\Pi_0), \Pi_0)$, and let A_0 be a y - $a(\Pi_0)$ -path in the resulting caterpillar (A_0 will be an initial subpath of A^i). If Π_1 lies in RL_-^i , then we can choose $A^i = A_0$, which is indeed an $\mathring{R}L_-^i$ -path in that case. If not, then we will go on recursively, trying in each step j to extend the already chosen initial subpath A_{j-1} of A^i , by attaching a path in $V(\Pi_j)$, where Π_j contains the endpoint of A_{j-1} , to reach a pi or ear $\Pi_{j+1} \preceq \Pi_j$. As we shall see, we will, sooner or later, land on RL_-^i .

Formally, for $j = 1, 2, \dots$ perform a step of the following kind. Suppose that Π_j, A_{j-1} have been defined. If a bond b of Π_j is not present in $G^{\#}$, that is, there is no edge in $G^{\#}$ between the endvertices of b , then metamorphose the larva $a(\Pi_j)\Pi_j b$, and let A_j be the concatenation of A_{j-1} with an A_{j-1} - $a(\Pi_j)$ -path in the resulting caterpillar — as we shall see, Π_j could not have been handled while constructing an A^k for some $k < i$. Let Π_{j+1} be the pi or ear that contains $a(\Pi_j)$ as an inner vertex (note that $a(\Pi_j) \neq y$, since no anchors are sent to an articulation point). If both bonds of Π_j are single edges, then there are two cases.

In the case that $y(\Pi_j)$ is incident with a double edge f on Π_j , delete f and metamorphose the larvae $a(\Pi_j)\Pi_j f$ and $b(\Pi_j)\Pi_j f$. Let W be the one of these larvae meeting A_{j-1} , and let A_j be the concatenation of A_{j-1} with a path in the caterpillar of W connecting A_{j-1} to the mouth s of W (note that $y \neq b(\Pi_j)$, because otherwise we would have chosen Π_j rather than Π_0). Let Π_{j+1} be the pi or ear containing s as an inner vertex (thus $\Pi_{j+1} \preceq \Pi_j$). A detour for f will be specified later.

In the case that $y(\Pi_j)$ is incident with no double edge on Π_j , metamorphose the larvae $a(\Pi_j)\Pi_j y(\Pi_j)$ and $b(\Pi_j)\Pi_j y(\Pi_j)$. Let A_j be the concatenation of A_{j-1} with an A_{j-1} - $a(\Pi_j)$ -path in the resulting caterpillars. Let Π_{j+1} be the pi or ear that contains $a(\Pi_j)$ as an inner vertex.

In all cases, if Π_{j+1} lies in RL_-^i we stop the recursion and let $A^i = A_j$, which is by construction an $\overset{\circ}{RL}^i$ -path with precisely one endpoint at y . We call it the *apophysis of RL^i* . If Π_{j+1} does not lie in RL_-^i , we proceed with the next step. Clearly $\Pi_{j+1} \preceq \Pi_0$, and furthermore $\Pi_0^i \preceq \Pi_{j+1}$, because otherwise the $G_{i-1}^{\#}$ -bridge in which Π_0 lies meets both y_0^i and $G_{i-1}^{\#}$, contradicting the choice of y_0^i . Since there are only finitely many pis or ears Π with $\Pi_0^i \preceq \Pi \preceq \Pi_0$, the procedure will stop after $k \in \omega$ steps, with Π_{k+1} lying in RL_-^i .

With a similar argument, we see that as promised above, Π_j could not have been handled while constructing an A^k for some $k < i$. For if A^k uses Π_j , then as A^k has to reach Π_0^k or $\Pi_1^k \preceq \Pi_0^k$, it has to go through Π_0^i (recall that Π_j lies in a $(G_{i-1}^{\#} \cup RL_-^i)$ -bridge that meets $y = y_0^i$, and thus has all feet in $\Pi_0^i \cup \Pi_1^i$). But then, Π_0^i would have been handled before beginning with the construction of A^i , and we would have proceeded to step $i + 1$ without ever trying to construct A^i . In particular, A^i, A^k are disjoint if $i \neq k$.

The following observation will be useful in Section 7.4.5 where we will “clean up” the articulation points.

Observation 4. *If A^i contains an edge f incident with an articulation point $y_i^k \neq y$, then either f lies in RL^k or it represents two edges that lie in RL^k .*

Indeed, if Observation 4 is false, then pick the least j such that A_j contains an edge f contradicting it. Since by construction all edges added to A_j in

step j either lie in Π_j or represent edges of Π_j , f is the last edge of A_j and its incident vertex in Π_{j+1} is an articulation point y_l^k . But then, A_j yields a path (after replacing representing edges with the edges they represent) in G^\sharp that lies in a $(G_{k-1}^\sharp \cup RL_-^k)$ -bridge and connects y_l^k to $RL^i \preceq RL^k$, contradicting the choice of y_l^k .

We now divide Case II into three subcases, depending on where the end-point $y' \neq y$ of A^i lies. In all cases, our aim is to split $\Pi_0^i \cup \Pi_1^i$ in a set of larvae \mathcal{W}_i , so that (in addition to Conditions 1 and 2) the following two conditions are satisfied (note that these conditions are also satisfied in Case I; Condition 4 is even satisfied by pis used by some apophysis):

Condition 3. *The union of A^i with the graph induced by $V(\Pi_1^i \cup \Pi_0^i - b(\Pi_0^i))$ after metamorphosing all larvae in \mathcal{W}_i is connected.*

Condition 4. *r_2^i, l_2^i lie in the same larva in \mathcal{W} .*

First we consider the case $y' \in \Pi_1^i - r_1^i$ (Figure 7.11). If e is double, then switch W_1^i . Now e is single and e' double; delete e' . Then metamorphose the trivial larva (r_1^i, e) , and the larva $l_0^i \Pi_0^i l_1^i \Pi_1^i e'$. Pick a detour $dt(e')$ for e' in the union of A^i with the resulting caterpillar. Handle Π_0^i like in Case I: if the multiedge $d = l_1^i l'$ of P_0^i incident with l_1^i is double, delete it and metamorphose the larva $r_0^i \Pi_0^i l'$; pick a detour $dt(d)$ for d in the resulting caterpillars and A^i . If d is single, and there is a double edge f on P_0^i (Figure 7.11), delete f and metamorphose the larva $l_1^i P_0^i f$ and the larva $r_0^i \Pi_0^i f$; pick a detour $dt(f)$ in the resulting caterpillars and A^i . If there is no double edge on P_0^i , let r' be the neighbour of r_1^i on P_0^i , metamorphose the larva $l_1^i P_0^i r'$ and completely metamorphose the larva $r_0^i \Pi_0^i r'$; a detour for the resulting double caterpillar can again be found in the resulting caterpillars and A^i .




Figure 7.11: Splitting into larvae: Case II, $y' \in \Pi_1^i - r_1^i$, and there is a double edge f on P_0^i .

In the case that $y' \in l_0^i \Pi_0^i l_1^i$, switch W_1^i if needed, so that e is single and e' double; delete e' . Then metamorphose the trivial larva (r_1^i, e) , and the larva $r_0^i \Pi_0^i l_1^i \Pi_1^i e'$ (Figure 7.12). Pick a detour $dt(e')$ for e' in the union of A^i with the resulting caterpillars. Then, if the first multiedge $h = l_1^i l''$ of $l_1^i \Pi_0^i y'$ is double, delete it and metamorphose the larva $l_0^i \Pi_0^i l''$; pick a detour $dt(h)$ for h in the resulting caterpillars and A^i . If h is single, and there is a double edge f on $l_1^i \Pi_0^i y'$, delete it and metamorphose the larva $l_1^i \Pi_0^i f$ and the larva $l_0^i \Pi_0^i f$; pick a detour $dt(f)$ in the resulting caterpillars and A^i . If there is no double edge on $l_1^i \Pi_0^i y'$, let z be the neighbour of y' on $l_1^i \Pi_0^i y'$, metamorphose the larva $l_1^i \Pi_0^i z$ and completely metamorphose the larva $l_0^i \Pi_0^i z$ (Figure 7.12); a detour for the resulting double caterpillar can again be found in the resulting caterpillars and A^i .




Figure 7.12: Splitting into larvae: Case II, $y' \in l_0^i \Pi_0^i l_1^i$, and no double edge on $l_1^i \Pi_0^i y'$.

Finally, if $y' \in r_0^i \Pi_0^i l_1^i$, switch W_1^i if needed, so that e is double and e' single; delete e . Metamorphose the larva $l_0^i \Pi_0^i l_1^i \Pi_1^i y$. If the multiedge $d = l_1^i l'$ of P_0^i incident with l_1^i is double, delete it and metamorphose the larva $r_0^i \Pi_0^i l'$; pick a detour $dt(d)$ for d in the resulting caterpillars and A^i . If d is single, and there is a double edge f on $l_1^i \Pi_0^i y'$, delete it and metamorphose the larva $l_1^i P_0^i f$ and the larva $r_0^i \Pi_0^i f$; pick a detour $dt(f)$ in the resulting caterpillars and A^i . If there is no double edge on $l_1^i \Pi_0^i y'$, let w be the neighbour of y' on $l_1^i \Pi_0^i y'$, metamorphose the larva $l_1^i P_0^i w$ and completely metamorphose the larva $r_0^i \Pi_0^i w$; a detour for the latter larva can again be found in the resulting caterpillars and A^i . A detour $dt(e)$ for e can always be found in the resulting caterpillars and A^i .

It is easy to confirm that the following is true:

Observation 5. *No detour specified in Case II meets any $pi \Pi \neq \Pi_0^i$ for which $\Pi \preceq \Pi_0^i$ holds.*

Now is the time to specify a detour $dt(d)$ for each edge d we deleted during the construction of A^i . It will suffice to construct paths D_1, D_2 each

connecting an endpoint of d to $RL_-^i \cup A^i$. Then, since D_1, D_2 can only meet RL_-^i in Π_0^i or Π_1^i by the construction of RL^i , we can, by Condition 3, find a path D with vertices in $V(\Pi_0^i \cup \Pi_1^i \cup A^i)$ connecting the endpoints of D_1, D_2 , and set $dt(d) = D_1 \cup D \cup D_2$.

Deleting d separated the pi or ear on which it lies in two subpaths Q_1, Q_2 , which have already been metamorphosed, and one of them, say Q_1 , meets A^i , so we can choose D_1 to be a d - A^i -path in the corresponding caterpillar. In order to choose D_2 , we imitate the procedure we used to construct A^i : we split the pi or ear on which Q_2 lands in one or two larvae, unless it has already been handled (that is, split in larvae), making the same distinction of cases as we did for Π_j while constructing A^i , and prolong our current path by a path in the new caterpillars that brings us a bit closer to RL_-^i (or A^i). We repeat until we meet $RL_-^i \cup A^i$.

While constructing $dt(d)$, we might delete other double edges. But then, we just repeat the procedure recursively to find detours for them as well. Since, easily, any deleted edge lies in a pi or ear Π for which $\Pi_0^i \preceq \Pi \preceq \Pi_0$ holds, this will happen only finitely often. Moreover, the following is true:

Observation 6. *If a detour for an edge deleted while constructing A^i meets some pi or arc Π , then $\Pi_0^i \preceq \Pi \preceq \Pi_0$ holds (Π_0 was defined while constructing A^i).*

The first phase is now completed, and we proceed to the second. Let $(\pi_i)_{i \in \mathbb{N}}$ be an enumeration of the pis that were not handled above, so that $i \leq j$ if $\pi_i \preceq \pi_j$. For $i = 1, 2, \dots$, if a bond b of $\Pi := \pi_i$ is not present in G^{\emptyset} , metamorphose the larva $a(\Pi)\Pi b$. If not and both multiedges on Π incident with $y := y(\Pi)$ are single, metamorphose the larvae $a(\Pi)\Pi y$ and $b(\Pi)\Pi y$. Otherwise, delete a double edge f incident with y , and metamorphose the larvae $a(\Pi)\Pi f$ and $b(\Pi)\Pi f$. Note that in this case, $\Pi = \Pi_l^k$ for some k and $l > 0$, and Π_{l-1}^k has already been handled. By Condition 4 and by the way the pis in this phase are handled, $a(\Pi)$ and $b(\Pi)$ lie in the same larva W of Π_{l-1}^k . Pick a detour $dt(f)$ for f in the union of the caterpillar of W with the caterpillars of the larvae of Π . Clearly, the following is true:

Observation 7. *$dt(f)$ does not meet any pi $\Pi \neq \Pi_{l-1}^k$ for which $\Pi \preceq \Pi_{l-1}^k$ holds.*

Having handled all pis, we go on to the ear decompositions. For every ear decomposition D with ears C_1, C_2, \dots, C_k , recursively for $i = k, k-1, \dots, 1$, if C_i has not been handled yet (while constructing some apophysis), then we want to split C_i into larvae, so that we can move from any vertex of C_i towards some RL_-^n , without using an edge incident with some y_l^j ; more

precisely, we will split C_i into larvae, metamorphose them, and perhaps make some shortcuts, so that after all changes have been made to C_i , the following condition is satisfied:

Condition 5. *For every $x \in V(C_i)$, there is a path that connects x to some π_i or ear $\Pi \preceq C_i, \Pi \neq C_i$, and contains no edge incident with some y_l^j .*

We consider two cases. For the first case, if $C_i \cap G^\emptyset$ does not meet any y_l^j , then we treat it similarly with a π_i in $(\pi_i)_{i \in \mathbb{N}}$: if a bond b of C_i is not present in G^\emptyset , we metamorphose the larva $a(C_i)C_i b$. If not and both multiedges on C_i incident with $y := y(C_i)$ are single, we metamorphose the larvae $a(C_i)C_i y$ and $b(C_i)C_i y$. Otherwise, we delete a double edge f incident with y , and metamorphose the larvae $a(C_i)C_i f$ and $b(C_i)C_i f$; a detour for f will be specified later. Clearly, Condition 5 is now satisfied.

In the second case, when $C_i \cap G^\emptyset$ meets y_l^j for some j, l , note that both bonds of C_i must be present in G^\emptyset , as by definition the anchor of C_i does not meet y_l^j . Now if both multiedges on C_i incident with $y := y(C_i)$ are single, metamorphose the larvae $a(C_i)C_i y$ and $b(C_i)C_i y$. Otherwise, as the bonds of C_i are single edges, and y is incident with a double edge, there is in $a(C_i)C_i y$ a vertex incident with a single as well as a double edge in C_i ; let u be the first vertex in $a(C_i)C_i y$ with that property. All vertices have even degree in the current simple multigraph; indeed, we started with the eulerian simple multigraph G^\emptyset , and the operations we have been performing (see list after Condition 2) preserve the parities of the vertex degrees. Thus u has an odd number of edges in some $RL_-^j \cup \bigcup_{n \leq i} C_n$ -bridge B in the current simple multigraph; clearly, all vertices of B lie in $\bigcup_{n > i} C_n$, so B is finite. Again since all vertices have even degree, in particular those in B , by the ‘‘hand-shaking’’ lemma B has at least one foot $v \neq u$ in $RL_-^j \cup \bigcup_{n \leq i} C_n$; let P be a u - v -path in B . We consider three subcases:

If $v \notin C_i$, then there is no double edge in $a(C_i)C_i u$ by the choice of u , so let u' be the neighbour of u on $a(C_i)C_i u$, metamorphose the larva $a(C_i)C_i u'$ and completely metamorphose the larva $y_l^j C_i u'$. We claim that Condition 5 is now satisfied for C_i . Indeed, $V(C_i)$ is divided in a caterpillar X and a double caterpillar Y , and if $x \in V(C_i) - \{a(C_i), b(C_i)\}$ lies in X , then there is an x - $a(C_i)$ -path in X , whereas if x lies in Y , then by the construction of a double caterpillar, either there is an x - u -path in Y avoiding y_l^j , which can be extended by P to an x - v -path, or there is an x - u' -path in Y avoiding y_l^j , which can be extended by a u' - $a(C_i)$ -path in X to an x - $a(C_i)$ -path.

If $v \in C_i$, and there is no double edge in $v C_i u$, then $u \in y_l^j C_i v$. Let u' be the neighbour of u on $v C_i u$, completely metamorphose the larva $y_l^j C_i u'$ and metamorphose the larva $a(C_i)C_i u'$ (even if $u' = v$). By a similar argument as in the previous subcase, we see that again Condition 5 is satisfied.

If $v \in C_i$, and there is a double edge f in vC_iu , delete f and metamorphose the larvae $W_1 := a(C_i)C_if$ and $W_2 := y_l^j C_if$. To see that Condition 5 is satisfied, note that if x is a vertex in W_2 , then there is in the caterpillar of W_2 a path connecting x to P that avoids y_l^j , and there is in the caterpillar of W_1 a path connecting the endpoint of P to $a(C_i)$.

In the last subcase, if in addition $v = y_l^j$ then let X be the caterpillar containing v , and shortcut the edges of P, X incident with y_l^j ; call the new edge a *shortcutting edge*. Note that this change does not affect the satisfaction of Condition 5 by the ears in $\bigcup_{n>i} C_n$; neither does it affect any apophysis by Observation 4. Moreover, we claim that the shortcutted edges did not lie in any detour. Indeed, there are two kinds of vulnerable detours: those defined while constructing A^j , and those defined while handling the ears of D . For the former, note that by the choice of Π_0 in the construction of A^j , we have $\Pi_0 \preceq C_i$ because C_i is a candidate for Π_0 , and by Observation 6 no detour of the first kind was affected. For the latter, note that we have not yet specified any detours for deleted edges in D , apart from those automatically specified when metamorphosing a larva. But if y_l^j lies in a larva $W = (s, P)$ in D , then it is easy to check that $y_l^j = s$ by construction, and since s has degree 1 in P and in the caterpillar of W , no such detour goes through y_l^j . This completes the proof of the claim.

We need to specify detours for the edges of D that we deleted and for the double caterpillars. For every deleted edge e (respectively double caterpillar X), pick paths P_1, P_2 in the new graph, each connecting a different endvertex of e (a vertex of a different component of X) to $V - V(\bigcup D)$, which exist by Condition 5. Let Π be the lowest pi with respect to \preceq that $\bigcup D$ sends a bond to, and let Π' be a pi for which $\Pi' \prec \Pi$ holds (unless $\Pi = \Pi_0^0$, in which case let $\Pi' = \Pi$). By Condition 3 and the way we handled the pis in the second phase, a path P_3 connecting the endpoints of P_1, P_2 can be chosen, that does not meet any pi lower than Π' with respect to \preceq . Let $dt(e)$ (respectively $dt(X)$) be the path $P_1 \cup P_2 \cup P_3$.

This completes the second phase. Denote the resulting simple multigraph by G^\triangleleft . Let $G_1 := (V, E(G^\triangleleft) \cup E(G^\emptyset))$. Easily, by Lemma 2.10, $|G_1| \cong |G^\emptyset|$. The set $\{dt(e) | e \in E(G_1) - E(G^\triangleleft)\}$ is thin (if $e \in E(G_1) - E(G^\triangleleft)$ is one of the parallel edges belonging to a double edge e' , then take $dt(e)$ to equal $dt(e')$ if only the latter has been defined), since each time we chose some $dt(e)$ we specified a pi Π_0 , such that no $\Pi' \preceq \Pi_0$ could meet $dt(e)$ (see Observations 3 and 5 to 7 and the relevant remark in the previous paragraph), and no pi can have been specified as Π_0 infinitely often. Thus, again by Lemma 2.10, $|G^\triangleleft| \cong |G_1| \cong |G^\emptyset|$.

By Condition 4 and by the way that the pis in the second phase were handled, we obtain:

Observation 8. $V(A^i \cup RL^i - b(\Pi_0^i))$ induces a connected subgraph of G^\triangleleft for every i .

(Where we assume that A^i is the empty graph if it has not been defined.)

7.4.5 Cleaning up the articulation points

Keeping to our plan, we now rid the articulation points of unwanted edges. For every $i, j \in \mathbb{N}$, let F be the set of edges incident with y_j^i that have an endvertex outside $V(RL^i \cup A^i)$. By the construction of G^\sharp and G^\triangleleft , every element of F is or represents a bond, and as double bonds were deleted while constructing G^\emptyset , there is no pair of parallel edges in F . Now let f_1, f_2, \dots, f_k be an enumeration of F , and for $l = 1, 2, \dots, \lfloor \frac{k}{2} \rfloor$, shortcut f_{2l-1} with f_{2l} . Call the new edges *shortcutting edges* (recall that we have already defined another kind of shortcutting edges in Section 7.4.4). We are left with a simple multigraph G^y , where each y_j^i is incident with at most one edge not in R^i ; indeed, even if A^i exists, $|F|$ is even in that case because of parity reasons.

Nothing needs to be done at articulation points of ears, because they do not have any unwanted edges by construction. Again, we claim that we didn't change the end topology.

Let $G_2 := (V, E(G^\triangleleft) \cup E(G^y))$. Applying Lemma 2.10 to G_2, G^\triangleleft , using as a detour $dt(e)$ for each edge e in $E(G_2) - E(G^\triangleleft)$ the two edges of G^\triangleleft shortcutting to give e , we prove that $|G_2| \cong |G^\triangleleft|$.

We want to specify a detour for each deleted edge and apply Lemma 2.10. For each edge $e = uv \in E(G_2) - E(G^y)$, either e is a bond, or it represents a bond of Π , where Π is either a Π_0^l for some l , or an ear. Let y_j^i be the articulation point where e was shortcutting, and suppose that $u = y_j^i$. Note that by Observation 4 (Section 7.4.4) no edge of an apophysis was shortcutting.

In the case that $\Pi = \Pi_0^l$ for some l , we have $RL^i \preceq RL^l$ by Observation 1, thus there is a finite sequence of rope-ladders R_1, R_2, \dots, R_k such that $RL^i = R_k \prec R_{k-1} \prec \dots \prec R_1 = RL^l$. Let P_0 be the trivial path v . For $j = 1, 2, \dots, k-1$, there is by Observation 8 a path P_j in G^\triangleleft connecting the last vertex of P_{j-1} (which lies in R_j) to the anchor a_j of R_j (which lies in R_{j+1}) such that all vertices of P_j other than a_j lie in \mathring{R}_j and its apophysis. Let $P = P_0 \cup P_1 \cup \dots \cup P_k$. We claim that P , which was defined as a path in G^\triangleleft , is also a path in G^y . Thus we need to prove that no edge of P was shortcutting. We only shortcutting edges that meet two rope-ladders, and any such edge in P either lies in an apophysis, and is thus not shortcutting as mentioned above, or is or represents an anchor, in which case it meets no articulation point by the definition of anchor. This proves our claim that P is a path in G^y ; let a be its endvertex in RL^i .

As a is a foot of a G_i^\sharp -bridge in G that also has the articulation point y_j^i as a foot, a lies in $\Pi_j^i \cup \Pi_{j+1}^i$ by the construction of G^\sharp . Thus, by the construction of G^\triangleleft , there is an a - u -path Q in G^\triangleleft containing only vertices of $\Pi_{j-1}^i, \Pi_j^i, \Pi_{j+1}^i$ and A^i and, easily, Q is also a path in G^y . Thus we may choose $dt(e) := P \cup Q$ as a detour for e . Call P the P -part of $dt(e)$ and call Q the Q -part of $dt(e)$.

In the case that Π is an ear, by recursively applying Condition 5, we obtain a v - RL_-^i -path containing no edge incident with a y_l^j . As in the first case, we can augment this path by a path containing only vertices of $\Pi_{j-1}^i, \Pi_j^i, \Pi_{j+1}^i$ and A^i to obtain a detour $dt(e)$.

We claim that the set $\{dt(e) | e \in E(G_2) - E(G^y)\}$ is thin. We have to show that for any edge f there are only finitely many edges e such that $dt(e)$ contains f . It is not hard to see that there can only be finitely many such e that are or represent bonds of ear decompositions. If there are infinitely many such e that are or represent bonds of rope-ladders, then either there are infinitely many e such that the P -part of $dt(e)$ contains f , or infinitely many e such that the Q -part of $dt(e)$ contains f . Again, it is not hard to see that the latter cannot be the case. To see that the former cannot be the case either, note that if the P -part of $dt(e)$ contains f , then e is incident with a vertex that lies in a π that is lower with respect to \preceq than the π containing both vertices of f . Clearly, there are only finitely many such π s, and as each of them contains finitely many vertices of finite degree, there can only be finitely many such e . This completes the proof that the set $\{dt(e) | e \in E(G_2) - E(G^y)\}$ is thin, thus by Lemma 2.10, $|G^y| \cong |G_2| \cong |G^\triangleleft|$.

We further claim that G^y is eulerian. Let $G_3 = (V, E(G^\emptyset) \cup E(G^y))$. Easily, by Lemma 2.10, $|G_3| \cong |G^\emptyset|$, and since $|G^y| \cong |G^\triangleleft| \cong |G^\emptyset| \cong |G^\emptyset|$, we have $|G^y| \cong |G_3|$. We know that G^\emptyset is eulerian, thus, by Lemma 7.2 and the definition of the cycle space, $E(G^\emptyset)$ is the sum of a thin family F of circuits in G^\emptyset . Since $|G^\emptyset| \cong |G_3|$ and $|G^\emptyset| \subseteq |G_3|$, every element of F is also a circuit in G_3 . Now let $T := E(G^y) \Delta E(G^\emptyset)$, where Δ denotes the symmetric difference. Clearly, T can be expressed as the sum of a thin set of finite cycles, since in order to get G^y from G^\emptyset we performed a number of operations, each of which consisted in either replacing a path of length 2 with an edge forming a triangle with the path, or deleting a double edge, or switching a window (see the list of allowed operations after Condition 2), and no edge participated in more than two such operations. But then, $E(G^y) = T \Delta E(G^\emptyset)$ holds, which means that $E(G^y)$ is the sum of the thin family $F \cup T$ of circuits in G_3 , thus an element of the cycle space of G_3 . By Lemma 2.1, $E(G^y)$ is a set of disjoint circuits in G_3 , and since $|G^y| \cong |G_3|$, these circuits are also circuits in G^y , proving that G^y is eulerian.

7.4.6 The hamiltonisation

By Theorem 7.3 we obtain an end-faithful Euler tour σ of G^y . Replace every shortcutting edge in σ by the two edges it shortcuts; formally, this is done by modifying σ on the interval of S^1 mapped to the shortcutting edge, so that this interval is mapped continuously and bijectively to the two shortcutted edges. Then, replace every representing edge in the resulting mapping by the two edges it represents, to obtain a mapping $\sigma' : S^1 \rightarrow G^{\circ\circ}$, where $G^{\circ\circ}$ is the simple multigraph resulting from G° after doubling all single edges; σ' is clearly end-faithful.

A *pass* (of σ') through some vertex x , is a trail $uex'e'v$ traversed by σ' . *Lifting* a pass $P = uex'e'v$ is the operation of replacing P in σ' with a u - v -edge if $u \neq v$ (formally, this is done by modifying σ' on the interval of S^1 mapped to P , so that this interval is mapped continuously and bijectively to the u - v -edge), or replacing P in σ' with the trivial trail u if $u = v$. As e, e' are edges of $G^{\circ\circ}$, uv is an edge of G^2 in the first case. Our plan is to perform some lifts so as to transform σ' into a Hamilton circle of G^2 , so we will first mark some passes for later lifting, then show that no two passes share an edge, and thus we can do lift them all at once without creating any edge not in G^2 .

For every $x \in V - \{y^*\}$, let i be the index of x in $P(W(x))$, and mark all passes of σ' through x that do not contain $e_i^-(W(x))$. Moreover, mark all passes of σ' through y^* that do not contain e^* (recall that e^* , the special edge in the assertion of Theorem 7.7, is an anchor, thus it has not been deleted). We claim that for every edge e traversed by σ' , at most one of the two passes that contain e was marked, which implies that no two passes share an edge.

In order to prove this claim, suppose that e is an edge with endvertices x, v and that the (unique) pass through x containing e has been marked. If $x = y^*$, then easily $e = e_1(W)$, where $W = W(v)$, thus the pass through $v = x_1(W)$ containing e has not been marked. If $x \neq y^*$, then let $W = W(x)$ and suppose that $x = x_i(W)$. Again we will show that the pass through v containing e has not been marked.

If e lies in $P(W)$, then $e \neq e_i^-$ because the pass through x containing e has been marked. Moreover, $e \neq e_{i+1}^+$, because if e_{i+1}^+ exists, then e_i^-, e_{i+1}^+ had been represented in G^{\triangleleft} , and thus e_{i+1}^+ lies in the pass through $x = x_i$ that contains e_i^- . If $e = e_i^+$, then by the same argument, it lies in the pass through x_{i-1} that contains e_{i-1}^- , which, according to our rules for marking, has not been marked. If $e = e_{i+1}^-$, again the pass through x_{i+1} that contains e cannot be marked, unless x_{i+1} is the pseudo-mouth of W ; but if $x_{i+1} = v$ is the pseudo-mouth of W , then e, e_i^- were represented in G^{\triangleleft} , so they both lie in the pass of σ' through $x = x_i$. But, according to our marking rules,

this pass cannot have been marked, contradicting our assumption.

If e does not lie in $P(W)$, let W' be the larva in \mathcal{W} in which e lies (it must lie in one). If x is the mouth of W' , then $v = x_1(W')$, $W(v) = W'$ and $e = e_1^-(W')$, thus the pass through v containing e has not been marked. If x is the pseudo-mouth of W' , then $x = x_k(W')$, where $k = |P(W')|$, $v = x_{k-1}(W')$ and $e = e_k^-(W')$. But $e_k^-(W'), e_{k-1}^-(W')$ were represented in G^\triangleleft , so they both lie in the pass through v , which, according to our marking rules, was not marked. The only case left, by Condition 1, is that x is an articulation point and it is the last vertex of both W, W' . In this case, both $e, e_i(W)$ are single edges by the construction of G^\triangleleft , and they are the only edges incident with x in G^y by the construction of G^y . But then, they both lie in the pass through x , contradicting the fact that this pass has been marked.

Thus our claim is proved, and so we can lift all marked passes at once without creating any edge not in G^2 . This transforms σ' to a mapping $\tau : S^1 \rightarrow |G^2|$. It is not hard to see that no pass of σ' through some vertex $v \neq y^*$ containing an edge incident with y^* could have been marked (see the beginning of the proof of our claim), and hence $\tau(S^1)$ contains e^* , and the other edge in $\tau(S^1)$ incident with y^* is also in $E(G)$.

By Lemma 2.10 we easily have $|G^2| \cong |G|$, and as $|G| \cong |G^\#|$ and, trivially, $|G^\#| \cong |G^{\emptyset\emptyset}|$, it follows that τ is continuous and end-faithful. Since for any vertex $v \in V$, all passes through v but for precisely one pass were marked and eventually lifted, τ traverses each vertex in V exactly once. In particular, τ does not contain any pair of parallel edges, and we can therefore replace each edge in τ that is parallel to an edge e in G with e , to obtain a Hamilton circle of G^2 . This completes the proof of Theorem 7.7, which implies Theorem 7.2.

The fact that the square of a 2-connected finite graph G is Hamiltonian connected ([32]), also generalises to locally finite graphs:

Corollary 2. *The square of a 2-connected locally finite graph G is Hamiltonian connected, that is, for each pair of vertices x, y of G , there is a homeomorphic image in $|G^2|$ of the unit interval with endpoints x, y .*

Proof. Add a new vertex y^* to G , join it to x, y with edges and apply Theorem 7.7. \square

7.5 Final remarks

In this chapter we generalised Fleischner's Theorem to locally finite graphs. What about generalising other sufficient conditions for the existence of a Hamilton cycle? In general, as in our case, it is a hard task, and it is not

clear why it should be possible. See for example [8], where Tutte's Theorem [30], that a finite 4-connected planar graph has a Hamilton cycle, is partly generalised. However, if instead of a Hamilton circle, we demand the existence of a closed topological path that traverses each vertex exactly once, but may traverse ends more than once, the task becomes much easier. Usually, one only has to apply the sufficient condition for finite graphs on a sequence of growing finite subgraphs of a given infinite graph G and use compactness, to obtain such a topological path in $|G|$. The difficult problem is how to guarantee injectivity at the ends. Here we used Theorem 7.3 to overcome this difficulty. A general approach suggests itself: try to reduce the existence of a Hamilton cycle in a finite graph, to the existence of some suitable Euler tour in some auxiliary graph, and then try to generalise the proof to the infinite case using Theorem 7.3.

The following easy corollary of Theorem 7.3 is perhaps an argument in favour of this approach:

Corollary 3. *If G is a locally finite eulerian graph, then its line graph $L(G)$ has a Hamilton circle.*

Proof. If R is a ray in G , then $E(R)$ is the vertex set of a ray $l(R)$ in $L(G)$. It is easy to confirm that the map

$$\begin{aligned} \pi : \Omega(G) &\rightarrow \Omega(L(G)) \\ \omega &\mapsto \omega' \ni l(R), R \in \omega \end{aligned}$$

is well defined, and it is a bijection.

Now let σ be an end-faithful Euler tour of G , that maps a closed interval on each vertex of G . Let $\sigma' : S^1 \rightarrow |L(G)|$ be a mapping defined as follows:

- σ' maps the preimage under σ of each edge $e \in E(G)$ to $e \in V(L(G))$;
- for each interval I of S^1 mapped by σ to a trail $xeye'w$, σ' maps the subinterval I' of I mapped to y , continuously and bijectively to the edge $ee' \in E(L(G))$;
- σ' maps the preimage under σ of each end $\omega \in \Omega(G)$ to $\pi(\omega)$.

Then “contract” in σ' each interval mapped to a vertex to a single point, to obtain the mapping $\tau : S^1 \rightarrow |L(G)|$. Since, in locally finite graphs, every finite vertex set is incident with finitely many edges, and every finite edge set is covered by a finite vertex set, $\Omega(G)$ and $\Omega(L(G))$ have the same topology. Thus τ is continuous and injective, and since S^1 is compact and $|L(G)|$ Hausdorff, a homeomorphism. Clearly, it traverses each vertex of $|L(G)|$ exactly once. \square

Using the finite version of Corollary 3, Zhan [33] proved that every finite 7-connected line graph is hamiltonian. In view of Corollary 3, a generalisation to locally finite graphs looks plausible:

Conjecture 7.1. *Every locally finite 7-connected line graph has a Hamilton circle.*

Bibliography

- [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** (2004), 235–256.
- [3] H. Bruhn and R. Diestel, *Duality in infinite graphs*, To appear in *Comb., Probab. Comput.*
- [4] H. Bruhn, R. Diestel, and M. Stein, *Cycle-cocycle partitions and faithful cycle covers for locally finite graphs*, *J. Graph Theory* **50** (2005), 150–161.
- [5] H. Bruhn and A. Georgakopoulos, *Linear algebra with infinite sums*, Preprint 2006.
- [6] H. Bruhn and M. Stein, *MacLane’s planarity criterion for locally finite graphs*, To appear in *J. Combin. Theory (Series B)*.
- [7] ———, *On end degrees and infinite circuits in locally finite graphs*, To appear in *Combinatorica*.
- [8] H. Bruhn and X. Yu, *Hamilton cycles in infinite planar graphs: see Chapter 6 of H. Bruhn, Infinite circuits in locally finite graphs, PhD thesis, Universität Hamburg, 2005.*
- [9] S.N. Burris, *Logic for mathematics and computer science*, Prentice Hall, 1998.
- [10] R. Diestel, *End spaces and spanning trees*, To appear in *J. Combin. Theory (Series B)*.
- [11] ———, *The cycle space of an infinite graph*, *Comb., Probab. Comput.* **14** (2005), 59–79.

- [12] ———, *Graph theory* (3rd edition), Springer-Verlag, 2005,
Electronic edition available at:
<http://www.math.uni-hamburg.de/home/diestel/books/graph.theory>.
- [13] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory (Series B) **87** (2003), 197–206.
- [14] ———, *On infinite cycles I*, Combinatorica **24** (2004), 68–89.
- [15] ———, *On infinite cycles II*, Combinatorica **24** (2004), 91–116.
- [16] ———, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.
- [17] H. Fleischner, *On spanning subgraphs of a connected bridgeless graph and their application to DT-graphs*, J. Combin. Theory (Series B) **16** (1974), 17–28.
- [18] ———, *The square of every two-connected graph is hamiltonian*, J. Combin. Theory (Series B) **16** (1974), 29–34.
- [19] H. Freudenthal, *Über die Enden topologischer Räume und Gruppen*, Math. Zeitschr. **33** (1931), 692–713.
- [20] A. Georgakopoulos, *Connected but not path-connected subspaces of infinite graphs*, To appear in *Combinatorica*.
- [21] ———, *Infinite hamilton cycles in squares of locally finite graphs*, Preprint 2006.
- [22] ———, *A short proof of Fleischner’s theorem*, Preprint 2006.
- [23] A. Georgakopoulos and P. Sprüssel, *Geodesic topological cycles in locally finite graphs*, Preprint 2006.
- [24] R. Halin, *A note on Menger’s theorem for infinite locally finite graphs*, Abh. Math. Sem. Univ. Hamburg **40** (1974), 111–114.
- [25] D.W. Hall and G.L. Spencer, *Elementary topology*, John Wiley, New York 1955.
- [26] P. Heinrich, *Abzählbar-unendliche Graphen mit HAMILTONscher dritter Potenz*, Math. Nachr. **82** (1978), 231–239.
- [27] D. König, *Theorie der endlichen und unendlichen graphen*, Akademische Verlagsgesellschaft, 1936.

- [28] M. Sekanina, *On an ordering of the set of vertices of a connected graph*, Publ. Fac. Sc. Univ. Brno **412** (1960), 137–142.
- [29] C. Thomassen, *Hamiltonian paths in squares of infinite locally finite blocks*, Annals of Discrete Mathematics **3** (1978), 269–277.
- [30] W.T. Tutte, *A theorem on planar graphs*, Trans. Am. Math. Soc. **82** (1956), 99–116.
- [31] ———, *How to draw a graph*, Proc. London Math. Soc. **13** (1963), 743–768.
- [32] S. Říha, *A new proof of the theorem by Fleischner*, J. Combin. Theory (Series B) **52** (1991), 117–123.
- [33] S. Zhan, *On hamiltonian line graphs and connectivity*, Discrete Mathematics **89** (1991), 89–95.

Angelos Georgakopoulos
Topological paths and cycles in infinite graphs
Hamburg 2007

Zusammenfassung

Das Hauptresultat dieser Dissertation ist die Verallgemeinerung auf lokal-endlichen Graphen des bekannten Satzes von Fleischner (Kapitel 7). Der Satz von Fleischner besagt dass das Quadrat jedes 2-zusammenhängenden endlichen Graphen Hamiltonsch ist. Diese Aussage wird hier für lokal endlichen Graphen bewiesen; die Definition vom Hamiltonkreis die dabei verwendet wird ist diejenige von Bruhn: ein (topologischer) Hamiltonkreis ist ein homöomorphes Bild von S^1 in der Freudenthal Kompaktifizierung $|G|$ des Graphen das alle Ecken enthält.

Ein Nebenresultat des entsprechenden Beweises ist ein kurzer Beweis des Satzes von Fleischner (Kapitel 7).

Ein weiteres Resultat dieser Dissertation ist dass die geodätische Kreise eines lokal-endlichen Graphen G , bezüglich einer Zuweisung von Längen zu den Kanten von G , den topologischen Zyklusraum $\mathcal{C}(G)$ von G erzeugen (Kapitel 6). Der topologische Zyklusraum wurde von Diestel und Kühn eingeführt, und hat die Verallgemeinerung von mehreren grundsätzlichen Eigenschaften des endlichen Zyklusraumes auf lokal-endlichen Graphen ermöglicht. Das benannte Resultat ist eine Verallgemeinerung dieser Art.

Desweiteren, es wurde durch Angabe eines Beispiels bewiesen, dass es einen lokal-endlichen Graphen G gibt, so dass $|G|$ einen Teilraum X besitzt, der topologisch zusammenhängend aber nicht wegzusammenhängend ist (Kapitel 4). Dies widerlegt eine Vermutung von Diestel.

Lebenslauf

Persönliche Daten

Name	Angelos Georgakopoulos
Geburtsdatum	26. April 1979
Geburtsort	Athen, Griechenland
Adresse	Henriettenstr. 6 20259 Hamburg Deutschland
Email	georgakopoulos@math.uni-hamburg.de

Universitäre Ausbildung

seit 10/2003	Universität Hamburg Promotionsstudium bei Reinhard Diestel
10/1996 – 02/2002	National Technical University of Athens Diplom Elektrik- und Computeringenieur

Sonstiges

10/2005 – 09/2006	Leistungsstipendium für ausländische Studierende der Universität Hamburg
6/1997	dritter Preis beim griechischen Mathematischen Wettbewerb für Schüler Einer von acht Schülern, die den dritten Preis bekommen haben