Infinite Hamilton Cycles in Squares of Locally Finite Graphs

Agelos Georgakopoulos

Abstract

We prove Diestel's conjecture that the square G^2 of a 2-connected locally finite graph G has a Hamilton circle, a homeomorphic copy of the complex unit circle S^1 in the Freudenthal compactification of G^2.

1 Introduction

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. A *Hamilton* cycle in a graph is a cycle containing all its vertices. Although Hamilton cycles are a central notion in graph theory and there is a vast literature about them, very few natural sufficient conditions are known for their existence. The following classical theorem of Fleischner [18] is perhaps the deepest known sufficiency result:

Theorem 1 (Fleischner, 1974)**.** *If G is a finite 2-connected graph, then its square G^2 has a Hamilton cycle.*

Thomassen [30] generalised Theorem 1 to locally finite graphs with one end:

Theorem 2 (Thomassen, 1978)**.** *If G is a 2-connected locally finite 1-ended graph, then G^2 contains both a Hamilton ray and a Hamilton double ray.*

A *double ray*, i.e. a two-way infinite path, is an infinite cycle in the algebraic sense of simplicial homology: if we orient it either way, the boundaries of its edges sum to zero. Thus in this sense Theorem 2 extends Theorem 1. Yet it is clear that no graph with more than two ends can contain a Hamilton double ray: since such a graph has a finite set of vertices separating it into more than two infinite components, no double ray can visit all its vertices. Hence there is no hope of generalising Thomassen's theorem, with cycles taken to be double rays, further to arbitrary 2-connected locally finite graphs.

However, things look better if we reinterpret Theorem 2 geometrically. In the case of a 1-ended graph G, a double ray is an infinite cycle also in the geometric sense of Diestel [10, 11]: its closure in the Freudenthal compactification $|G|$ of G is a topological *circle*, a subspace homeomorphic to the circle S^1. A *Hamilton circle* of G, then, is a circle in $|G|$ that contains every vertex of G.

In three seminal papers [14, 15, 16], Diestel and Kühn established that this geometric notion of a cycle can serve as a basis for a homology of locally finite graphs. This has since been shown by various authors [3, 4, 5, 6, 23] to outperform both the simplicial homology of G (in which all cycles are finite) and the so-called *open homology* alluded to above (in which all double rays are cycles): in the context of graph homology, at least, topological circles appear to be the 'right' analogue of the cycles in finite graphs.

Motivated by this development, Diestel [12] suggested an ambitious programme to use topological paths and circles as a basis also for a translation of mainstream 'extremal' finite graph theory to locally finite graphs. As a benchmark test for the feasibility of such a programme, he conjectured that the notion of a Hamilton circle should make it possible to unify Fleischner's and Thomassen's theorems into a general theorem for arbitrary locally finite graphs: that the square of any 2-connected locally finite graph has a Hamilton circle [10].

Our aim in this paper is to prove Diestel's conjecture:

Theorem 3. *If G is a locally finite 2-connected graph, then G^2 has a Hamilton circle.*

One of the ideas used for the proof of Theorem 3 led to a short proof of Theorem 1, which will be published separately [20].

As an intermediate step, we obtain a result which may be of independent interest. A *topological Euler tour* of G is a continuous map $\sigma : S^1 \to |G|$ that traverses every edge of G exactly once. Topological Euler tours are known to exist when expected, e.g. when every vertex and every end of G has even degree [6, 14]. A topological Euler tour is *injective at ends* if it traverses every end of G exactly once. As a lemma for the proof of Theorem 3, we shall prove the following:

Theorem 4. *If a locally finite multigraph has a topological Euler tour, then it also has one that is injective at ends.*

Theorem 4 might help generalise other sufficient conditions for the existence of Hamilton cycles in finite graphs to Hamilton circles in locally finite graphs; see Section 10 for details.

We shall also generalise to locally finite graphs the well-known fact, proved by Karaganis [26] and Sekanina [29], that the third power of any connected finite graph has a Hamilton cycle:

Theorem 5. *If G is a connected locally finite graph, then G^3 has a Hamilton circle.*

It is a well-known conjecture (apparently first stated in [28], see [2] for more) that every finite connected Cayley graph has a Hamilton cycle. Although it is not true that every locally finite connected Cayley graph has a Hamilton circle (see Section 9), Theorem 5 (or Theorem 3) implies the following:

Corollary 6. *Every finitely generated group Γ has a finite generator set S such that the Cayley graph of Γ with respect to S has a Hamilton circle.*

This paper is structured as follows. After providing the definitions and some basic results required in our proofs (Sections 2 and 3), we prove Theorem 5 in Section 4. The proof of Theorem 3 is sketched in Section 5 and completed in Sections 6 and 7. (Theorem 4 is proved in Section 6.) Section 8 offers some conjectures about Hamilton circles in infinite graphs that are not locally finite. Corollary 6 motivates some further problems, which we discuss in Section 9. We wind up in Section 10 with some concluding remarks.

2 Definitions

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

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

For any $v \in V$, let $G[v]_i$ be the subgraph of G induced by the vertices at distance at most i from v.

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

A *path* is a graph $P = (V(P), E(P))$ of the form $V(P) = \{x_0, x_1, \dots, x_k\}$, $E(P) = \{x_0x_1, x_1x_2, \dots x_{k-1}x_k\}$, where the x_i are all distinct. If $e \in E(P)$ and $x, y \in V(P)$, then xPe is the shortest subpath of P connecting x to an endvertex of e, xPy is the subpath of P connecting x to y, $\mathring{x}P\mathring{y} = xPy - \{x, y\}$, etc.

A *walk* in G is a non-empty alternating sequence $v_0e_0v_1e_1 \dots e_{k-1}v_k$ of vertices and edges in G such that $e_i = \{v_i, v_{i+1}\}$ for all $i < k$. A *trail* in G is a walk in which no edge appears more than once.

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 $\widehat{C}$ 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 $\widehat{C}$ is an *H-bridge* in G (in the literature an edge in $E(G) - E(H)$ with both endvertices in H is usually also a bridge, but in this paper bridges always contain more than one edge). Its *feet* are the vertices in $V(\widehat{C}) - V(C)$. An *H-path* in G is a path having precisely its endvertices (but no edge) in common with H.

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

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.

A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*. A *tail of* the ray R is a final subpath of R. Two rays R, L in G are *equivalent* if no finite set of vertices separates them; 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 ends of G by $\Omega = \Omega(G)$. A ray belonging to the end ω is an *ω-ray*.

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

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

and let $\Omega(S,\omega)$ denote the set of all ends of G with a ray in C. As our basis of open neighbourhoods of ω we now take all sets of the form

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

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

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

A *circle* in $|G|$ is the image of a homeomorphism from S^1, the unit circle in $\mathbb{R}^2$, to $|G|$. A *Hamilton circle of* G is a circle that contains every vertex of G (and hence, also every end, as it is closed). An *a*rc in $|G|$ is a homeomorphic image of the real interval $[0,1]$ in $|G|$.

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

A *topological Euler tour of* G, or *Euler tour* for short, is a continuous map $\sigma : S^1 \to |G|$ such that every inner point of an edge of G is the image of exactly one point of S^1 (thus, every edge is traversed exactly once, and in a "straight" manner). Call G *eulerian* if it has a topological Euler tour. A map $\sigma : S^1 \to |G|$ is *injective at ends* if every end in $\Omega(G)$ has exactly one preimage under σ.

3 Basic Facts

3.1 Infinite cycles and paths

The following two lemmas are perhaps the most fundamental facts about the cycle space of an infinite graph. Both can be found in [11, Theorem 8.5.8]. Let G be an arbitrary connected locally finite multigraph fixed throughout this section (the following results have been proved for simple graphs only, but they can be easily generalised to multigraphs.

Lemma 1. *Every element of* $\mathcal{C}(G)$ *is a disjoint union of circuits.*

Lemma 2. *Let* $F \subseteq E(G)$. *Then* $F \in \mathcal{C}(G)$ *if and only if* F *meets every finite cut in an even number of edges.*

The next lemma comes from [14, Theorem 7.2].

Lemma 3. *The following three assertions are equivalent:*

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

A continuous map from the real unit interval $[0,1]$ to a topological space X is a *(topological) path* in X. The following lemma can be found in [21]. It will be used in Section 4.

Lemma 4. *A topological path that connects some vertex or end of a basic open neighbourhood U of an end $\omega \in \Omega(G)$, to a vertex or end outside U, must traverse some edge xy with $x \in U, y \notin U$.*

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 [11, 8.2.2].

Lemma 5. *If U is an infinite set of vertices in G, then G contains a comb with all teeth in U.*

On two occasions in this paper we will make use of a proof technique which is very common in infinite graph theory and is referred to as compactness, see [11] for an introduction. Usually, the simplest way to use compactness is applying Theorem 7 below. In order to state it, suppose that we have specified a set $\mathcal{V}$, whose elements we will call the *logical variables*, and define a *propositional formula* with variables in $\mathcal{U}$, where $\mathcal{U}$ is a finite subset of $\mathcal{V}$, to be a subset P of the set $\{0,1\}^{\mathcal{U}}$ of functions $f : \mathcal{U} \to \{0,1\}$; intuitively, a logical variable represents a choice between two simple, complementary, facts — encoded by the two elements of $\{0,1\}$, which elements we call the *truth-values* — like for example the existence or not of an edge between two given vertices, and a propositional formula expresses the truth or not of some fact based on the logical variables, for example the existence or not of a cycle on a given finite vertex set: given an assignment of truth-values $g : \mathcal{V} \to \{0,1\}$, we say that g *satisfies* the propositional formula $P \subseteq \{0,1\}^{\mathcal{U}}$ if the restriction of g to $\mathcal{U}$ is an element of P. We will also say that g satisfies a set K of propositional formulas if it satisfies every element of this set. Finally, we say that K is *satisfiable* if there is a $g : \mathcal{V} \to \{0,1\}$ that satisfies K.

Theorem 7. *Let K be an infinite set of propositional formulas, every finite subset of which is satisfiable. Then K is satisfiable.*

Theorem 7 appears often in the literature, e.g. in [8], although the definitions of the concepts involved are usually phrased in different terminology. For the convenience of the interested reader we prove Theorem 7 here using a standard argument, as found e.g. in [11, Theorem 8.1.3].

Proof of Theorem 7. Consider the product space $\Pi := \{0,1\}^{\mathcal{V}}$ of $|\mathcal{V}|$ copies of the set $\{0,1\}$ endowed with the discrete topology (where $\mathcal{V}$ is the set of logical variables). By Tychonoff's theorem, this is a compact space. For every propositional formula P in K the set A_P of elements of Π that satisfy P is closed, as P depends only on finitely many components of Π. By our assumption the sets A_P have the finite intersection property, thus by the compactness of Π there is a point $p \in \Pi$ in their overall intersection $\bigcap_{P \in K} A_P$. It is straightforward to check that p satisfies K. □

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

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 situation is 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 $H \subseteq G$ note that any two equivalent rays in H are also equivalent rays in G, which allows us to define a mapping π_{HG} by

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

Lemma 6. *Let H, G be locally finite connected multigraphs such that $H \subseteq G$, $V(H) = V(G)$, and for any two rays R, L in H, if $R \approx_G L$ then $R \approx_H L$. 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 5 to obtain a comb in H with teeth in $V(R)$. The spine of this comb is a ray L in H such that $L \approx_G R$. Thus its end is mapped to ω by π_{HG}.

Since $H \subseteq G$, if S is a finite subset of $V(G)$ and $\omega \in \Omega(H)$ then $C_H(S, \omega)$ is a subgraph of $C_G(S, \omega)$, from which it follows easily that π_{HG} is continuous. Moreover, $\Omega(H)$ is compact, because it is closed in $|H|$ and $|H|$ is compact (see [11, Proposition 8.5.1]). 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 $\Omega(G)$. □

Lemma 7. *Let H, G be locally finite connected multigraphs such that $H \subseteq G$, $V(H) = V(G)$, and for any two rays R, L in H, if $R \approx_G L$ then $R \approx_H L$. 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 $\hat{\pi}_{HG} : V(H) \cup \Omega(H) \to V(G) \cup \Omega(G)$ that maps every end $\omega \in \Omega(H)$ to $\pi_{HG}(\omega)$, and every vertex in $V(H)$ to itself. Easily by Lemma 6, $\hat{\pi}_{HG}$ is bijective and continuous. Moreover, $V(H) \cup \Omega(H)$ is closed, thus compact, so like in the proof of Lemma 6, $\hat{\pi}_{HG}$ is a homeomorphism between $V(H) \cup \Omega(H)$ and $V(G) \cup \Omega(G)$, from which the assertion easily follows. □

For any two connected multigraphs G, H such that $V(G) = V(H)$, we will write $|H| \approxeq |G|$ if there is a homeomorphism $\pi : \Omega(H) \to \Omega(G)$, 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 $H \subseteq G$ are fixed 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 8. *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 $|H| \approxeq |G|$.*

Proof. Clearly, H is connected. Pick any two rays R, L in H, such that $R \approx_G L$. By Lemmas 6 and 7, it suffices to show that $R \approx_H L$.

Since $R \approx_G L$, there is an infinite set $\mathcal{P}$ of disjoint R-L–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-L–path in P'. The set of all these paths $\{dt(P)|P \in \mathcal{P}\}$ is clearly thin, from which it easily follows that $R \approx_H L$. □

4 The Third Power of a Locally Finite Graph is Hamiltonian

Karaganis [26] and Sekanina [29] have proved that the third power of a connected finite graph is hamiltonian. Extensions of this fact to infinite graphs have been achieved by Sekanina [29], who showed that the third power of a connected, locally finite, 1-ended graph has a spanning ray, and by Heinrich [25], who specified a class of non-locally-finite graphs whose third power has a spanning ray. With Theorem 5, which we prove in this section, we generalise to locally finite graphs with any number of ends.

Proof of Theorem 5. 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 [11, Theorem 8.2.4]), and let $T_i = T[v]_i$.

We will prove the assertion using Theorem 7. To this end, define for each edge $e \in E(T^3)$ a logical variable ℓ_e; the truth-values of ℓ_e will encode the presence or not of e. Let $\mathcal{V} := \{\ell_e \mid e \in E(T^3)\}$ 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. Finally, 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 $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2 \cup \mathcal{P}_3$.

In order to meet the condition of Theorem 7 we have to show that 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 $\mathcal{P}'$. This is indeed true: choosing i so large that T_i^3 contains all vertices and all finite cuts that come up in $\mathcal{P}'$ and applying the following lemma to T_i, yields a Hamilton cycle of T_i^3 which encodes such an assignment:

Lemma 9. *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 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. □

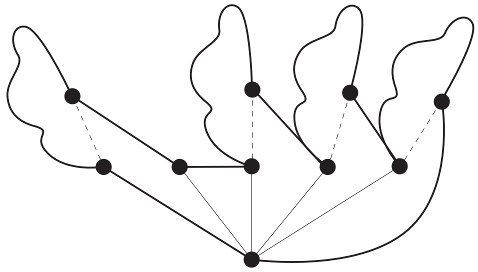

Figure 1: Using the induction hypothesis to construct a Hamilton cycle of T. The wavy curves represent Hamilton cycles of the v-branches supplied by the induction hypothesis, and for each such Hamilton cycle H, $e(H)$ is represented by a dashed line. The thick cycle represents H.

So by Theorem 7, 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 and the formulas in $\mathcal{P}_3$ we obtain $F \in \mathcal{C}(T)$, thus by Lemma 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 components B of $T - T_i$ such that $B \cap C \neq \emptyset$, and let $X = V(C) \cup \bigcup_{B \in \mathcal{B}} V(B)$ (note that any such B is a branch of T). 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 4 and Lemma 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 8 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 $|T^3| \approxeq |T|$, and similarly $|G^3| \approxeq |G|$ (these sets of detours are thin because the graphs are locally finite). Easily by Lemma 7, $|T| \approxeq |G|$, thus H is also a Hamilton circle of G^3. □

5 Outline of the Proof of Theorem 3

Before giving an outline of the proof of Theorem 3, let me compare it with the proof of Theorem 2 and a proof of Theorem 1 given by Říha ([32] or [11]) which is shorter than its original proof (in fact, Říha proves a stronger assertion than that of Theorem 1, see Section 7.1). The descriptions that follow are approximate, omitting much information not needed for the comparison.

Říha's proof uses induction; he finds a special cycle C, and then applies the induction hypothesis to every component of $G - C$, to obtain Hamilton cycles yielding 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 into a Hamilton cycle of G^2 by replacing some subtrails of length two of J with edges of G^2 having the same endvertices; we call this process the *hamiltonisation* of J.

Thomassen follows a similar plan in his proof of Theorem 2 (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 the finite theorem 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 Říha'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 2.

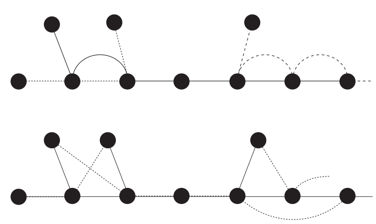

Figure 2: An example of a local hamiltonisation. In the upper figure, the restriction of the Euler tour on the segment (horizontal path) is indicated; it consists of three paths. In the figure beneath, these paths have been transformed into disjoint paths in the square of the graph that span all vertices (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 into 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 injective at ends. So we have to ask, which eulerian graphs admit an Euler tour that is injective at ends. The answer is given by Theorem 4: 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 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 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 could still continue with the plan of finding an Euler tour and transforming it to a Hamilton circle H of G', but even if this worked H would not be a Hamilton circle of G: it would traverse ω twice.

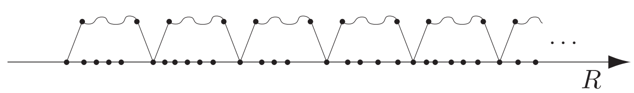

Figure 3: Performing all shortcuts between bonds would create a new end.

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 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 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 4. Thus 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 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 6, which shows the auxiliary graph corresponding to the graph of Figure 5. Note that the problematic set of paths in Figure 5 could not result from an Euler tour of a graph containing the graph in Figure 6. The idea of imposing such

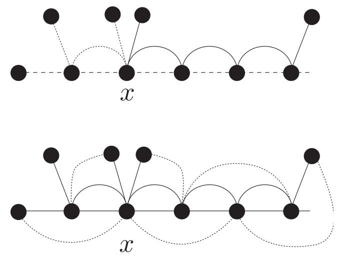

Figure 4: A difficult case: the lower figure shows the only possible hamiltonisation of the trails shown in the upper figure.

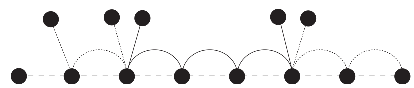

Figure 5: A case where no hamiltonisation is possible.

constraints on the Euler tour is used in [20] to obtain a short proof of Theorem 1.

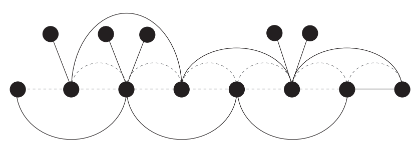

Figure 6: Applying constraints for an Euler tour of the graph of Figure 5. The dashed lines are edges of the original graph not in the auxiliary graph, while the continuous curved lines are edges of the auxiliary graph not in the original one.

The proof of Theorem 3 is structured as follows. We start by constructing the "scaffolding", that is, the analogue of R in Thomassen's proof, in Section 7.2. It consists of a set of ladder-like structures like the one shown in Figure 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 7.3. In Section 7.4, we will split $G^{\emptyset}$ in segments, called *larvae*, as follows. We consider each Π shaped subpath P of a rope-ladder like the thick path in Figure 7, called a *pi*, consisting of the two subpaths of the horizontal rays between two consecutive "rungs", and the rung on their right, and distinguish three cases. If one of the endedges of P became a double edge in $G^{\emptyset}$ — we have made sure that at most one did — we delete it, and consider the rest of P (or rather, the multigraph that replaced P in $G^{\emptyset}$) as a larva. If not, then we look at a special vertex in P carefully chosen while constructing the scaffolding, denoted by y^i_j and called an *articulation point*, and if one of the multiedges of P incident with the articulation point is double we delete it, and consider the two remaining subpaths of P as larvae. If both multiedges incident with the articulation point are single, we consider the two maximal subpaths of P ending at the articulation

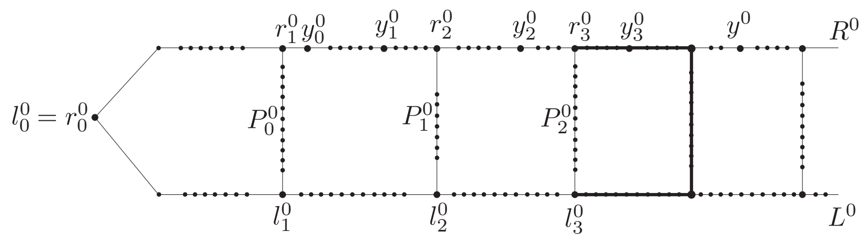

Figure 7: A rope-ladder. The horizontal paths are equivalent rays; usually their first vertices do not coincide as they do here.

point as larvae (Figure 8). The first pi of each rope-ladder, however, does not follow these rules, which is the reason for the anomaly regarding y_0^0 in Figure 7 (the articulation point corresponding to the first pi lies in the second one, which contains two articulation points, while each subsequent pi contains one articulation point). An ear decomposition is treated in a similar way. In all cases, we make sure that the first multiedge of every larva is a single edge.

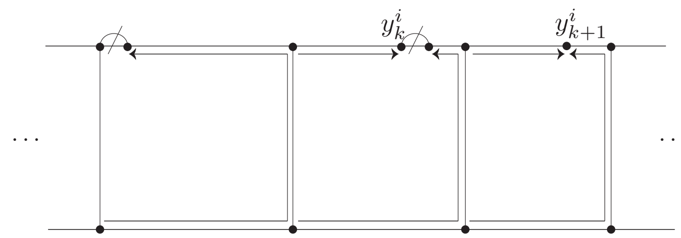

Figure 8: Splitting the graph into larvae. Every arrow indicates a larva.

Having divided the whole graph into larvae, we impose the aforementioned constraints on the Euler tour (in the same section). These constraints are so effective, that no shortcuts like the ones in the proofs of Říha and 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 (as in Figure 2). If larvae intersect otherwise however, there could be a conflict between their hamiltonisations. As shown in Figure 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 7.5, 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 7.6, pick the Euler tour that is injective at ends. 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 7.6 that it is possible to hamiltonise it to obtain a Hamilton circle of G^2.

Summing up, the proof of Theorem 3 consists of the following steps:

1. constructing the scaffolding;
2. making it eulerian;
3. splitting it into larvae;
4. imposing constraints on the Euler tour;
5. cleaning up the articulation points;
6. picking an Euler tour and hamiltonising it.

6 General Results

6.1 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.

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 10. *For every graph G and every countable end $\omega \in \Omega(G)$, if G has a set K of $k \in \mathbb{N}$ pairwise disjoint ω-rays, then it also has a set K' of k pairwise 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 [11, 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 after deleting R_0 the R_i do not have to be equivalent. However, it is easy to find a finite set $S \subset V$ such that any two tails of elements of K that lie in the same component of $G - R_0 - S$ are equivalent, and we can even choose S so that each R_j leaves S only once, because otherwise we can add an initial subpath of R_j to S. Applying the induction hypothesis to every component of $G - R_0 - S$ that contains 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. Then, let R'_0 be a ray in $G - \bigcup\{R'_1, R'_2, \dots, R'_{k-1}\}$ 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 $L \in \omega$, $L \cap \bigcup K' = \emptyset$, and let $\mathcal{P}$ be a set of infinitely many disjoint L-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, L 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, L is equivalent with some R'_i in $G - R_0 - S$ and thus meets some R'_j; in both cases the definition of L is contradicted, so our claim is true. □

Lemma 11. *If G is locally finite, $\omega \in \Omega$, and K is a set of ω-rays devouring ω in G, then every component of $G - \bigcup K$ sends finitely many edges to K.*

Proof. If such a component sends infinitely many edges to K then, easily, by Lemma 5 it contains a comb whose spine is equivalent with the rays in K, contradicting the assumption that K meets every ω-ray. □

It would be interesting to decide whether Lemma 10 remains true for infinite K:

Problem 1. *Let G be a graph, ω a countable end of G, and K an infinite set of pairwise disjoint ω-rays. Prove that there is a set K' of pairwise disjoint ω-rays that devours ω such that the set of starting vertices of rays in K' equals the set of starting vertices of rays in K.*

6.2 End-faithful topological Euler tours

In this section we prove Theorem 4.

Proof of Theorem 4. By Lemma 3 every finite cut of G is even. Then, G has a finite cycle C, because otherwise every edge would form a cut. Let $\sigma_0 : S^1 \to C$ 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 a topological Euler tour σ of G that is injective at ends. After each step i, we will have defined a finite set of edges F_i, which will be the union of a set of disjoint finite circuits, and a continuous surjection $\sigma_i : S^1 \to \bar{F}_i$ 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 in $E(G) - F_i$, 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 a vertex v incident with both D and F_{i-1} such that the distance between v and e_j is minimal, put v in S_i and let I_v be any of the closed intervals of S^1 mapped to v by σ_{i-1}. As F_{i-1} is a union of disjoint finite circuits, any finite cut in $G - F_{i-1}$ is even, since this was true for G and any finite cycle meets a finite cut in an even number of edges. Thus, every edge of D is contained in a finite cycle in D. Now choose a finite cycle C_D in D incident with v. 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 those 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 $F \subset F_j$ for j large enough, so denote by D, D' the components of $G - F_j$ that contain rays of ω, ω' respectively. But if x is mapped on a point p by σ_{j+1}, then for all steps succeeding $j+1$, x will be mapped on a point belonging to the component of $G - F_j$ that contains p. Thus $(\sigma_i(x))_{i\in\mathbb{N}}$ cannot meet both D, D' for $i > j$, a contradiction that proves the claim.

So we may define

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

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 for any $i \in \mathbb{N}$ is, clearly, 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. By the choice of v it traverses every edge, so it is 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 our claim is true and thus σ is injective at ends. $\square$

An *open Euler tour* of a locally finite multigraph G is a continuous mapping from the real unit interval $[0,1]$ to $|G|$ that traverses every edge of G exactly once. By the next corollary, Theorem 4 can be extended to open Euler tours; this will be used in Section 7.6 to obtain a corollary of Theorem 3, but not for the proof of Theorem 3 itself.

Corollary 8. *Let G be a locally finite multigraph, let $x, y \in V \cup \Omega$, and suppose that a finite cut in G is odd if and only if it separates x from y. Then, G has an open Euler tour with endpoints x, y that is injective at ends.*

Proof (sketch). We will only treat the case when x is a vertex and y an end; the other cases are similar. Let R be a path in G starting at x and devouring y, which exists by Lemma 10. By Lemma 11, every component of $G - R$ sends finitely many edges to R. For each such component C, let $\tilde{C}$ be the multigraph resulting from G after contracting $V - C$ to a single vertex v_C. It is easy to check that every finite cut in $\tilde{C}$ is even, so applying Theorem 4 we obtain an Euler tour

σ_C of $\tilde{C}$ that is injective at ends. We now use this fact to construct an auxiliary multigraph from G as follows. For every component C of $G - R$, decontracting v_C divides σ_C into a finite number of arcs with endpoints in $\widehat{C} \cap R$; remove C from G, and for any such arc P add an edge e_P, called an *arc-edge*, joining the endvertices of P. Doing so for all components C yields a new multigraph G' having only one end — that of R. By a result of Erdős et al. [17], a locally finite connected 1-ended graph H has a 1-way infinite trail starting at x and containing all edges of H if x is the only vertex of odd degree in H. Easily, every vertex of G' has the degree it had in G, thus x is the only vertex of odd degree in G' and we can apply this result. Writing the resulting 1-way infinite trail as a mapping $\sigma : [0,1] \to |G'|$, and replacing in σ each arc-edge e_P with P, yields the required open Euler tour.

□

7 Proof of Theorem 3

7.1 A stronger assertion

Říha's [32] proof of Theorem 1 mentioned in Section 5 proves in fact a slightly stronger assertion:

Theorem 9. *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 that contains e^* and a y^*-edge $e' \in E(G)$ with $e' \neq e^*$.*

Rather than Theorem 1, we will generalise this stronger assertion:

Theorem 10. *Let $G = (V, E)$ be an infinite 2-connected locally finite graph, let $y^* \in V$ and let $e^* = y^*x^* \in E$. Then, G^2 has a Hamilton circle that contains e^* and a y^*-edge $e' \in E$ with $e' \neq e^*$.*

7.2 Constructing the scaffolding

In this section we construct the "scaffolding" $G^\sharp$ mentioned in Section 5. The scaffolding will be made of two ingredients: *rope-ladders* and *ear decompositions*. Let us see 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 connected subgraph of B spanning $V(B - H)$ and containing some vertices of $B \cap H$ that consists of a sequence $C_1, C_2, \ldots, C_n$ 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 $|C_i| \geq 3$).

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 5) 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 contracting $G-D$ to a vertex z. If we wish to yield 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 12. *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 in B 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 through 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 through D 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 $i = 2, 3, \ldots$, suppose that $C_1, C_2, \ldots, C_{i-1}$ 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 $B-H$ not contained in $\bigcup_{j<i} C_j$, 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 $C_1, C_2, \ldots, C_{i-1}$ 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) 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 an articulation point lie in its pi, but instead we will, as explained in Section 5, 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 13. *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 13, there are two disjoint ω-rays starting at y^* and x^* (recall that y^*, x^* are the special vertices in the assertion of Theorem 10), and by Lemma 10 there is a pair $\{R^0, L\}$ of disjoint rays starting at y^*, x^* respectively that devours ω. Let $L^0 = y^*x^*L$, and let $r^0_0, l^0_0 = y^*$. Choose a sequence $(y^0_j)_{j\in\mathbb{N}}$ of vertices of R^0, and a sequence $(P^0_j)_{j\in\mathbb{N}}$ of pairwise disjoint R^0-L^0 paths, P^0_j having the endpoints r^0_{j+1}, l^0_{j+1}, so that y^0_0 is the first vertex on R^0 after r^0_1, and for each $j > 0$ the following conditions are satisfied (see Figure 7):

- y^0_j lies on $\mathring{y}^0_{j-1}R^0$;
- r^0_{j+1} lies in $\mathring{y}^0_j R^0 \mathring{y}^0_{j+1}$, and l^0_{j+1} lies in $\mathring{l}^0_j L^0$;
- Every $(R^0 \cup L^0)$–bridge that has y^0_{j-1} as a foot, has all other feet in $\mathring{r}^0_{j-1}R^0\mathring{r}^0_{j+1} \cup \mathring{l}^0_{j-1}L^0\mathring{l}^0_{j+1} - y^0_0$

(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^0_j and P^0_j in the order $P^0_0, y^0_1, P^0_1, y^0_2, P^0_2, \ldots$: recall that by Lemma 11 every $(R^0 \cup L^0)$–bridge has only finitely many feet, so each time we want to choose a new y^0_j or P^0_j, we just have to go far enough along R^0 and L^0.

Let RL^0 be the subgraph of G consisting of $RL^0_- := R^0 \cup L^0 \cup \{P^0_j | j \in \mathbb{N}\}$ and the ears of a fixed ear decomposition of every finite RL^0_-–bridge, which exists by Lemma 12. Let $\mathring{R}L^0 = RL^0$ and let $G^\sharp_0 = RL^0$.

The construction of $G^\sharp_0$ was the first step in an infinite procedure the aim of which is to define $G^\sharp$. Each step i of this procedure will be similar to the construction of $G^\sharp_0$: we will choose rays R^i, L^i in $G - G^\sharp_{i-1}$, and add them together with some R^i-L^i–paths and some ear decompositions to $G^\sharp_{i-1}$ to obtain $G^\sharp_i$. The endpoints of R^i, L^i will be distinct vertices of $G^\sharp_{i-1}$.

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 $G - G^\sharp_{i-1}$ containing x_j, where j is the smallest index so that $x_j \notin G^\sharp_{i-1}$; if no such j exists, then stop the procedure and set $G^\sharp = G^\sharp_{i-1}$. If the path Q_j has not been defined yet, then let it be any x_j-$\mathring{R}L^l$–path in $\widehat{C_i}$, 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 $G^\sharp_{i-1}$, and $w = w(i)$ the vertex after v on Q_j (thus $w \in G^\sharp_{i-1}$).

Intuitively, we want to have x_j in $G^\sharp_i$, but this might be impossible if x_j is "far" from $G^\sharp_{i-1}$, in which case we just try to make sure that $G^\sharp_i$ is closer to x_j than $G^\sharp_{i-1}$ 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 $G^\sharp_{i-1}$ at $\mathring{R}L^l$ is needed in order to guarantee that $G^\sharp$ 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 Section 4, 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 [11]). 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. *There are disjoint rays $R^i \approx L^i$ in $\widehat{C_i}$ so that: the first vertex of R^i is w and the first vertex of L^i lies in $G^\sharp_{i-1}$, the pair $\{R^i, L^i\}$ devours some end of G, and 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 13 and Lemma 10 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 $G^\sharp_i$). By Lemma 11, 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 $G^\sharp_{i-1}$-$(R' \cup L^*)$–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 $G^\sharp_{i-1} \cup R \cup L$ 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 $G^\sharp_{i-1}$, 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 $R^i =: r^i_0 R^i, L^i =: l^i_0 L^i$ having been chosen as in the Claim, pick a sequence $(y^i_j)_{j\in\mathbb{N}}$ of vertices of R^i, and a sequence $(P^i_j)_{j\in\mathbb{N}}$ of pairwise disjoint R^i-L^i paths in C_i, P^i_j having the endpoints r^i_{j+1}, l^i_{j+1}, so that y^i_0 is the first vertex on R^i after the endpoint of P^i_0, and for each $j > 0$ the following conditions are satisfied:

- y^i_j lies on $\mathring{y}^i_{j-1} R^i$;
- r^i_{j+1} lies in $\mathring{y}^i_j R^i \mathring{y}^i_{j+1}$, and l^i_{j+1} lies in $\mathring{l}^i_j L^i$;

- Every $(G^\sharp_{i-1}\cup R^i\cup L^i)$–bridge in G that has y^i_{j-1} as a foot, has all other feet in $\mathring{r}^i_{j-1}R^i\mathring{r}^i_{j+1}\cup \mathring{l}^i_{j-1}L^i\mathring{l}^i_{j+1}-y^i_0$

Such a choice is possible because by Lemma 11 every $(G^\sharp_{i-1}\cup R^i\cup L^i)$–bridge in G has finitely many feet, and there are only finitely many $(G^\sharp_{i-1}\cup R^i\cup L^i)$–bridges in G with feet on both $G^\sharp_{i-1}$ and $R^i\cup L^i$ (again, we choose the y^i_j and P^i_j in the order $P^i_0, y^i_1, P^i_1, y^i_2, P^i_2, \ldots$).

Let RL^i be the graph consisting of $RL^i_- := R^i\cup L^i\cup\{P^i_j | j\in\mathbb{N}\}$ 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 $\mathring{R}L^i = RL^i - \{r^i_0, l^i_0\}$. Recall that one of R^i, L^i contains an edge incident with w. Call this edge the *anchor* of RL^i, unless $w = y^k_j$ for some j,k, in which case let the other edge of $R^i\cup L^i$ incident with $G^\sharp_{i-1}$ 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^i_j, 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^i_j, P^i_{j-1} a *window* of RL^i, and denote it by W^i_j. Moreover, let Π^i_0 denote the path $r^i_0R^iP^i_0L^il^i_0$, and for any $j\geq 1$, let $\Pi^i_j = W^i_j - \Pi^i_{j-1}$. For every $i,j\in\mathbb{N}$, call Π^i_j a *pi*, let $y(\Pi^i_j) = y^i_j$, and call y^i_j an *articulation point*. The *bonds* of a pi are its endedges. The *bonds* of W^i_j are the bonds of Π^i_j and the *bonds* of RL^i are the bonds of Π^i_0 (that is, the endedges of RL^i). 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 $\mathring{R}L^j$, then $RL^j\preceq RL^i$.*

For suppose that RL^i sends a bond to $\mathring{R}L^j$ 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^i_0 be the anchor of RL^i. For every Π^i_j 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^i_j 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 $\Pi = \Pi^i_j$

and $\Pi' = \Pi^i_{j+1}$ for some i, j, or $\Pi' = \Pi^i_0$ 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 (consider y^* to be an inner vertex of Π^0_0). Let $\preceq$ be the reflexive transitive closure of $\prec$. Clearly, $\preceq$ is a partial order.

Define $G^\sharp_i$ as the union of $G^\sharp_{i-1}$ with RL^i and an ear decomposition of every finite $(G^\sharp_{i-1} \cup RL^i)$–bridge.

We can now define $G^\sharp := \bigcup_{i\in\mathbb{N}} G^\sharp_i$. 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 definition of x_j, any vertex $v \in V$ will at some step n either lie in $G^\sharp_n$, thus also in $G^\sharp$, or be chosen as x_j. 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, $v(i) \in G^\sharp_i$. Thus, at most $|Q_j|$ steps after step n (when the path Q_j was defined) x_j will lie in $G^\sharp_i$, which implies that $V(G^\sharp) = V$.

Our next aim is to prove that $|G^\sharp| \approxeq |G|$, and we will do so using Lemma 7. 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 $G^\sharp_j$ separates a tail of Q from a tail of T in $G^\sharp$ (just choose j large enough so that $G^\sharp_j$ 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 $G - G^\sharp_j$ 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-$\mathring{R}$–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-$\mathring{R}$–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 $G^\sharp_j$ separates Q from T in $G^\sharp$.

Thus no such rays Q, T exist and by Lemma 7, $|G| \approxeq |G^\sharp|$.

7.3 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 $G^\between$, 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 7. In order to meet its condition, we will show that:

Claim. *For every $i \in \mathbb{N}$ there is an eulerian simple multigraph $G^\between_i$ on V, so that any two vertices are neighbours in $G^\between_i$ 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 $G^\between_i$.*

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 the parity of vertex degrees.

In order to prove the Claim, begin by doubling all edges of $G^\sharp$. Then, for every ear decomposition $C_1, C_2, \ldots, C_k$ meeting $G^\sharp[y^*]_i$, recursively, for $j = k, k-1, \ldots, 0$, 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 $j = l, l-1, \ldots, 0$, 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 $G^\sharp_{j-1}$ that has the same endvertices as Π_0^j and contains no edge of an ear decomposition (for $j = 0$ switch Π_0^0 if its anchor is double). After the end of this recursion, switch every window whose anchor is a double edge and lies in $G^\sharp[y^*]_i$.

Let $G_i^\between$ be the resulting simple multigraph. Note that $G_i^\between$ 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, $G_i^\between$ is eulerian by Lemma 3. It is easy to see that $G_i^\between$ satisfies all conditions of the Claim. $\square$

In order to apply Theorem 7, 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.

Our last Claim implies that every finite set of these propositional formulas is satisfiable, so by Theorem 7 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 $G^\between$ which is eulerian (by Lemma 3), and in which all anchors of $G^\sharp$ form single edges.

Let $G^{\not\between}$ be the simple multigraph resulting from $G^\between$ after deleting each double edge that has the same endvertices as a bond in $G^\sharp$. Obviously, $|G^\between| \approxeq |G^\sharp|$ holds, and we claim that furthermore $|G^{\not\between}| \approxeq |G^\between|$. In order to prove this assertion, we will specify a thin set of detours for the deleted edges and apply Lemma 8.

If $e = pq$ is a deleted bond of a rope-ladder RL^i, let RL^j be the rope-ladder with the least index so that e meets $\mathring{RL}^j$, and suppose that p lies in $\mathring{RL}^j$ and q in $\mathring{RL}^i$. We claim that there is a p-q–path $dt(e)$ in $G^{\not\between}$ that satisfies the following conditions:

(i). $dt(e)$ is contained in the union of the 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 $G^\sharp - G^\sharp_j$ that contains $\mathring{RL}^i$.

To prove this, note that as each pi lost at most one bond and no other edges, $\mathring{RL}^l \cap G^{\not\between}$ 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 $RL^l \cap G^{\not\between}$ that contains 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 $G^{\not\between} \cap G^\sharp_i$ that is contained in the union of the 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 $RL^j \cap G^{\not\between}$ 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 $G^{\not\between}$ contained in the union of D with the rope-ladders R such that $RL^j \preceq R \preceq RL^i$ which avoids the pis of RL^j below the first one that meets the $(G^\sharp_{j-1} \cup RL^j_-)$-bridge in $G^\sharp$ 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 $\{dt(e) | e \in E(G^{\between}) - E(G^{\not\between})\}$ 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 $f \in \Pi^l_m$ 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 $dt(e)$ containins 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 $G^\sharp - G^\sharp_l$ that send an edge to some Π^l_k with $k \leq m$. Since the graph is locally finite, there are only finitely many such components, and by Lemma 11 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), $\mathring{R}$ and $\mathring{R}L^l$ lie in the same component C of $G^\sharp - G^\sharp_k$, and as e has to be one of the edges between C and $G^\sharp_k$, which again by Lemma 11 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 by the definition of $dt(e)$ any such ear decomposition D must lie in a $(G^\sharp_{l-1} \cup RL^l_-)$-bridge in $G^\sharp$ that has feet in both $G^\sharp_{l-1} \cup \bigcup_{k \leq m} \Pi^i_k$ and RL^l_-, and by the construction of $G^\sharp$ there are only finitely many such bridges; indeed, any such bridge lies in the $G^\sharp_{l-1}$-bridge in G in which RL^l lies, and this bridge has finitely many feet on $G^\sharp_{l-1}$ by Lemma 11. Again, every $(G^\sharp_{l-1} \cup RL^l_-)$-bridge in $G^\sharp$ sends finitely many edges to $(G^\sharp_{l-1} \cup RL^l_-)$ by Lemma 11, and as D must send an edge to $G^\sharp_{l-1} \cup \bigcup_{k \leq m} \Pi^l_k$ 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 $\{dt(e) | e \in E(G^{\between}) - E(G^{\not\between})\}$ is thin, which by Lemma 8 implies that $|G^{\not\between}| \approxeq |G^{\between}| \approxeq |G^\sharp|$.

7.4 Splitting into finite multigraphs

7.4.1 Larvae and caterpillars

While constructing $G^\sharp$ we defined many terms like pi, window, rope-ladder, etc. that were subgraphs of $G^\sharp$. We will use those names and symbols for $G^{\between}$ as well, but now we will mean the simple multigraphs in $G^{\between}$ that replaced the subgraphs of $G^\sharp$ that used to bear these names and symbols; thus when referring to $G^{\between}$, we will use Π^i_j to denote the subgraph of $G^{\between}$ spanned by the multiedges whose endvertices were joined by an edge of Π^i_j in $G^\sharp$, a bond (respectively anchor) is a multiedge whose endvertices where joined by a bond (resp. anchor) in $G^\sharp$, and so on. Moreover, xy denotes from now on the multiedge with endvertices x, y.

According to our plan, as stated in Section 5, 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 obtained from a path after replacing some of its edges with double edges — in $G^{\between}$, s is one of its endvertices, called its *mouth*, and the multiedge of P incident with s is a single edge (if it exists). 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 $e^-_i, e^\cap_i$, 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 for the larva (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 6: *metamorphosing* the larva $W =: (s, P)$, is the operation of replacing, in P and in $G^{\between}$, the edges $e^\cap_{j+1}, e^-_j$, for every j such that e_{j+1} is a double edge, with an $x_{j-1}x_{j+1}$ edge f_j (Figure 9). The edge f_j is called a *representing* edge and it *represents* the edges $e^\cap_{j+1}, e^-_j$. Note that $e^\cap_{j+1}, e^-_j, f_j$ 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 $dt(e) \subseteq X$ for e is automatically specified. These detours will be used after we are done metamorphosing larvae in order to show, using Lemma 8, that the end-topology did not change.

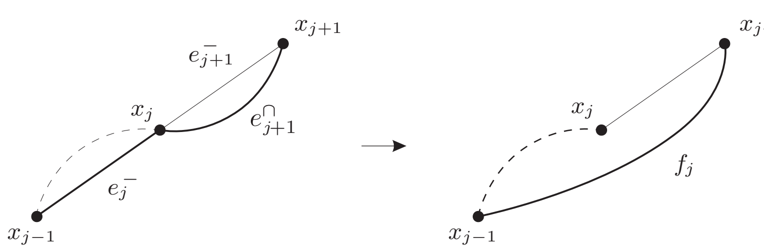

Figure 9: Replacing $e^\cap_{i+1}$ and e^-_i with a new edge f_j.

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 *two-headed caterpillar of* W is the graph X resulting from P after completely metamorphosing W. A two-headed 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 two-headed 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, a 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 for each edge e deleted while completely metamorphosing W to get X, a detour $dt(e)$ for e in $X \cup dt(X)$ is automatically specified.

We now divide the graph into larvae, and either metamorphose or completely metamorphose each of them. (According to the sketch of the proof in Section 5, we first split the graph into larvae and then impose the constrains on the Euler tour, but in fact these two steps will be performed simultaneously, the constrains being imposed by metamorphosing or completely metamorphosing the larvae.) Formally, we will specify a set of edge-disjoint larvae $\mathcal{W}$ so that $G^{\not\fish} = \bigcup_{W \in \mathcal{W}} P(W)$, and the following conditions are satisfied:

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

- x *is the mouth of* W *or* W'*;*
- x *is the pseudo-mouth of* W *or* W'*; or*
- x *is an articulation point, both* W, W' *end at* x*, 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* $x \in V - y^*$*, there is an element* $W(x)$ *of* $\mathcal{W}$ *containing* x *so that* x *is neither the mouth nor the pseudo-mouth of* $W(x)$ *(by Condition 1, there is at most one* $W \in \mathcal{W}$ *with this property, unless* x *is an articulation point; if there are more than one then we just pick one and call it* $W(x)$*).*

In the rest of this section we will construct a simple multigraph $G^{\triangleleft}$ on V by performing operations of the following kinds on $G^{\not\fish}$:

- replacing two incident edges e, f with an edge forming a triangle with e, f;
- switching a window (we defined switching in the beginning of Section 7.3);
- adding a double edge from $G^{\fish} - G^{\not\fish}$;
- 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 8 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 5, and in particular in Figure 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 $V(\mathring{R}L_-^i)$ 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.

7.4.2 The first phase

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 into larvae, in some previous step, or if one of its bonds e is not present in $G^{\not\triangleleft}$ — that is, if there is no edge in $G^{\not\triangleleft}$ connecting the endpoints of e — proceed with the next step. Otherwise, if zw is a bond of Π_1^i that is not present in $G^{\not\triangleleft}$, 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 $y := y_0^i$ in 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 $l_0^i \Pi_0^i l_1^i \Pi_1^i y$ (Figure 10; recall that by the definition of $G^{\not\triangleleft}$, any bond present in it is a single edge). Then, 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 union of the three resulting caterpillars (that is, the caterpillars of (r_1^i, e), $l_0^i \Pi_0^i l_1^i \Pi_1^i y$ and $r_0^i \Pi_0^i l'$). 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 union of 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 $r_0^i \Pi_0^i r'$ (Figure 11); a detour for the two-headed caterpillar X of $r_0^i \Pi_0^i r'$ can be found in the union of 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.*

Recall that whenever we metamorphose a larva we assume that for each deleted edge a detour is chosen that lies in the resulting caterpillar. Observation 3 refers to these detours as well as the ones explicitely specified above. 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.

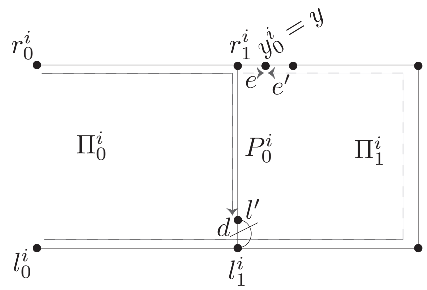

Figure 10: Splitting into larvae: Case I, and d is double. The dashed lines indicate larvae, and arrows show away from the mouth.

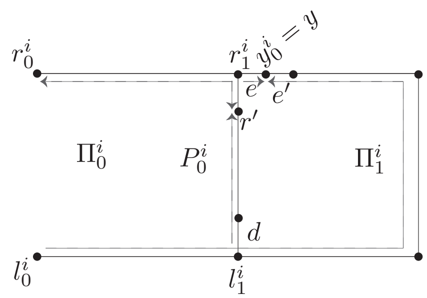

Figure 11: Splitting into larvae: Case I, and no double edge on P_0^i. The line with arrows at both ends indicates a larva that will be completely metamorphosed.

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, so as to obtain an $\mathring{R}L^i_-$–path A^i with one endpoint at y in the union of the resulting caterpillars, which path 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 b cannot be the anchor of Π_0, 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 will be a pi or ear containing 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 $j = 1, 2, \ldots$ 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 $G^{\emptyset}$, that is, there is no edge in $G^{\emptyset}$ 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 present in $G^{\not\sharp}$, then they are both single edges and we distinguish 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 two larvae that meets 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; thus $s \neq y$). 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 (much) 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 union of 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 $\mathring{R}L^i_-$–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 $\Pi^i_0 \preceq \Pi_{j+1}$, because otherwise the $G^\sharp_i$–bridge in which Π_0 lies meets both y^i_0 and $G^\sharp_{i-1}$, contradicting the choice of y^i_0. Since there are only finitely many pis or ears Π with $\Pi^i_0 \preceq \Pi \preceq \Pi_0$, 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^k_0 or $\Pi^k_1 \preceq \Pi^i_0$, it has to go through Π^i_0 (recall that Π_j lies in a $(G^\sharp_{i-1} \cup RL^i_-)$-bridge that meets $y = y^i_0$, and thus has all feet in $\Pi^i_0 \cup \Pi^i_1$). But then, Π^i_0 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. This implies in particular that A^i, A^k are disjoint if $i \neq k$.

The following observation will be useful in Section 7.5 where we will "clean up" the articulation points.

Observation 4. *If A^i contains an edge f incident with an articulation point $y^k_l \neq y$, 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^k_l. But then, A_j yields a path (after replacing representing edges with the edges they represent) in $G^\sharp$ that lies in a $(G^\sharp_{k-1} \cup RL^k_-)$-bridge and connects y^k_l to $RL^i \preceq RL^k$, contradicting the choice of y^k_l.

We now divide Case II into three subcases, depending on where the endpoint $y' \neq y$ of A^i lies. In all cases, our aim is to split $\Pi^i_0 \cup \Pi^i_1$ 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 3. *The union of A^i with the graph induced by $V(\Pi^i_1 \cup \Pi^i_0 - b(\Pi^i_0))$ after metamorphosing all larvae in $\mathcal{W}_i$ is connected.*

Condition 4. *r^i_2, l^i_2 lie in the same larva in $\mathcal{W}$. If some pi Π^k_j was handled while constructing A^i, then r^k_{j+1}, l^k_{j+1} lie in the same larva in $\mathcal{W}$.*

First we consider the case $y' \in \Pi_1^i - r_1^i$ (Figure 12). If e is double, then switch the window 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. Next, 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 union of the resulting caterpillars and A^i. If d is single, and there is a double edge f on P_0^i (Figure 12), 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 union of 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 two-headed caterpillar can again be found in the resulting caterpillars and A^i.

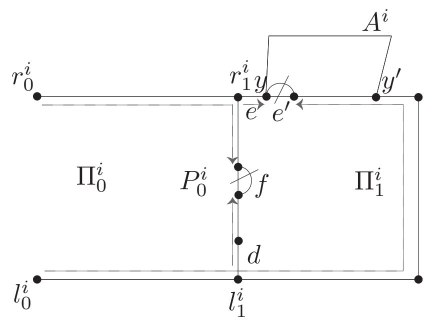

Figure 12: Splitting into larvae: Case II, $y' \in \Pi_1^i - r_1^i$, and there is a double edge f on P_0^i.

In the case that $y' \in \mathring{l}_0^i\Pi_0^i\mathring{l}_1^i$, switch W_1^i if needed so as to make e 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 13). Pick a detour $dt(e')$ for e' in the union of 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 $l_1^i\Pi_0^i z$ (unless $l_1^i = z$) and completely metamorphose the larva $l_0^i\Pi_0^i z$ (Figure 13); a detour for the resulting two-headed caterpillar can again be found in the resulting caterpillars and A^i.

Finally, if $y' \in \mathring{r}_0^i\Pi_0^i\mathring{l}_1^i$, switch W_1^i if needed so as to make e 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 z be the neighbour of y' on $l_1^i\Pi_0^i y'$, metamorphose the larva $l_1^i P_0^i z$ and completely metamorphose the larva $r_0^i\Pi_0^i z$; 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.

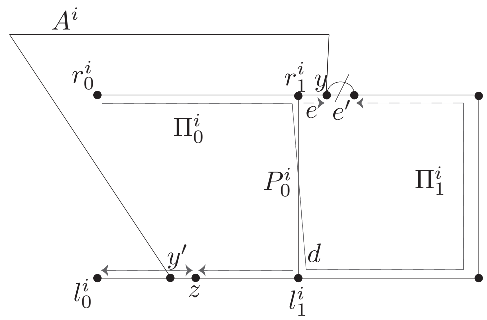

Figure 13: Splitting into larvae: Case II, $y' \in \mathring{l}^i_0 \Pi^i_0 \mathring{l}^i_1$, and no double edge on $l^i_1 \Pi^i_0 y'$.

It is easy to confirm that the following is true:

Observation 5. *No detour specified in Case II meets any pi* $\Pi \neq \Pi^i_0$ *for which* $\Pi \preceq \Pi^i_0$ *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 a distinct endpoint of d to $RL^i_- \cup A^i$. Then, since D_1, D_2 can only meet RL^i_- in Π^i_0 or Π^i_1 by the construction of RL^i, we can, by Condition 3, find a path D with vertices in $V(\Pi^i_0 \cup \Pi^i_1 \cup A^i)$ 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 $RL^i_- \cup A^i$; we will meet it sooner or later, by the argument that showed that A^i has to meet RL^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 $\Pi^i_0 \preceq \Pi \preceq \Pi_0$ holds (Π_0 was defined while constructing A^i), this will happen only finitely often. All these detours are chosen in already metamorphosed parts of the graph, and are thus immune to further changes. Note that Condition 4 still holds. 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* $\Pi^i_0 \preceq \Pi \preceq \Pi_0$ *holds.*

The first phase is now completed.

7.4.3 The second phase

We proceed to the second phase. 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 $i = 1, 2, \ldots$, if a bond b of

$\Pi := \pi_i$ is not present in $G^{\between}$, 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, $\Pi = \Pi^k_l$ for some k and $l > 0$, and Π^k_{l-1} 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^k_{l-1}. 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 $\Pi \neq \Pi^k_{l-1}$ for which $\Pi \preceq \Pi^k_{l-1}$ holds.*

Having handled all pis, we go on to the ear decompositions. For every ear decomposition D with ears $C_1, C_2, \ldots, C_k$, recursively for $i = k, k-1, \ldots, 1$, 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^j_l; 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 pi or ear $\Pi \preceq C_i, \Pi \neq C_i$, and contains no edge incident with any y^j_l.*

We consider two cases. For the first case, if $C_i \cap G^{\between}$ does not meet any y^j_l, then we treat it similarly with a pi in $(\pi_i)_{i \in \mathbb{N}}$: if a bond b of C_i is not present in $G^{\between}$, 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 $C_i \cap G^{\between}$ meets y^j_l for some j, l, note that both bonds of C_i must be present in $G^{\between}$, as by definition the anchor of C_i does not meet y^j_l. 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 $G^{\between}$, 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 outside C_i. By the definition of the articulation points, and as the ear decomposition D meets y^j_l, all these edges lie in $RL^j_- \cup \bigcup_{n \leq i} C_n$–bridges. 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 $y^j_l C_i u'$ (we will specify a detour for the two-headed caterpillar later). We claim that Condition 5 is now satisfied for

C_i. Indeed, $V(C_i)$ is divided in a caterpillar X and a two-headed 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 two-headed 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 vC_iu, then it follows from the definition of u that $u \in y_l^j C_i v$. Let u' be the neighbour of u on vC_iu, completely metamorphose the larva $y_l^j C_i u'$ and metamorphose the larva $a(C_i)C_iu'$ (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 $W_2 := y_l^j C_i f$. 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. Thus our claim is true.

We need to specify detours for the edges of D that we deleted and for the two-headed caterpillars. For every deleted edge e (respectively two-headed caterpillar X), pick paths P_1, P_2 in the new graph, each connecting a distinct endvertex of e (a vertex of a distinct 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 in the new graph 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 $G_1 := (V, E(G^{\triangleleft}) \cup E(G^{\between}))$. Easily, by Lemma 8, $|G_1| \approxeq |G^{\between}|$. 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 $dt(e')$ in case only the latter has been defined), since each time we chose some $dt(e)$ we specified a pi Π_0, such that no pi $\Pi' \prec \Pi_0$ 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 8, $|G^{\triangleleft}| \approxeq |G_1| \approxeq |G^{\between}|$.

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.)

7.5 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 in $G^{\triangleleft}$ 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 $G^{\not{\emptyset}}$, 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 7.4.3). 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 8 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 $G^{\triangleleft}$ shortcutted to give e, we prove that $|G_2| \approxeq |G^{\triangleleft}|$.

We want to specify a detour for each deleted edge and apply Lemma 8. 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 Π_0^l for some l, or an ear. Let y_j^i be the articulation point where e was shortcutted, and suppose that $u = y_j^i$. Note that by Observation 4 (Section 7.4.2) no edge of an apophysis was shortcutted.

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 by induction) to the anchor a_j of R_j (which lies in R_{j+1} by the definition of $\prec$) 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 shortcutted. We only shortcutted edges that meet two rope-ladders or a rope-ladder and an ear decomposition, and any such edge in P either lies in an apophysis, and is thus not shortcutted 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, clearly, 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^j_l. As in the first case, we can augment this path by a path containing only vertices of $\Pi^i_{j-1}, \Pi^i_j, \Pi^i_{j+1}$ 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 pi that is lower with respect to $\preceq$ than the pi containing both vertices of f. Clearly, there are only finitely many such pis, 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 8, $|G^y| \approxeq |G_2| \approxeq |G^{\lhd}|$.

We further claim that G^y is eulerian. Let $G_3 = (V, E(G^{\between}) \cup E(G^y))$. Easily, by Lemma 8, $|G_3| \approxeq |G^{\between}|$, and since $|G^y| \approxeq |G^{\lhd}| \approxeq |G^{\not\between}| \approxeq |G^{\between}|$, we have $|G^y| \approxeq |G_3|$. We know that $G^{\between}$ is eulerian, thus, by Lemma 3 and the definition of the cycle space, $E(G^{\between})$ is the sum of a thin family F of circuits in $G^{\between}$. Since $|G^{\between}| \approxeq |G_3|$ and $|G^{\between}| \subseteq |G_3|$, every element of F is also a circuit in G_3. Now let $T := E(G^y) \triangle E(G^{\between})$, where $\triangle$ 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 $G^{\between}$ 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, $E(G^y) = T \triangle E(G^{\between})$ 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 1, $E(G^y)$ is a set of disjoint circuits in G_3, and since $|G^y| \approxeq |G_3|$, these circuits are also circuits in G^y, thus G^y is eulerian by Lemma 3.

7.6 The hamiltonisation

By Theorem 4 we obtain an Euler tour σ of G^y that is injective at ends. 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 $\sigma' : S^1 \to G^{\between\between}$, where $G^{\between\between}$ is the simple multigraph resulting from $G^{\between}$ after doubling all single edges; σ' is clearly injective at ends.

A *pass (of σ')* through some vertex x, is a trail $uexe'v$ traversed by σ'. *Lifting* a pass $P = uexe'v$ is the operation of replacing P in σ' with a u-v–edge if $u \neq v$, or replacing P in σ' with the trivial trail u if $u = v$ (again, this is done by modifying σ' on the interval of S^1 mapped to P, so that this interval is either mapped continuously and bijectively to the u-v–edge or mapped to u). As e, e' are edges of $G^{\between\between}$, 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 $P(W(x))$ (see Condition 2 for the definition of $W(x)$), 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 10, 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 marked 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, $e \neq e_{i+1}^{\cap}$, because if $e_{i+1}^{\cap}$ exists, then $e_i^-, e_{i+1}^{\cap}$ had been represented in $G^{\vartriangleleft}$, and thus $e_{i+1}^{\cap}$ lies in the pass through $x = x_i$ that contains e_i^-. If $e = e_i^{\cap}$, 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^- where represented in $G^{\vartriangleleft}$, 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. If x is an articulation point and it is the last vertex of both W, W', then both $e, e_i(W)$ are single edges by the construction of $G^{\vartriangleleft}$, and they are the only edges incident with x in G^y by Condition 1 and the construction of G^y. But then, they both lie in the pass through x, contradicting the fact that this pass has been marked. If x is the mouth of W', then $v = x_1(W')$, $e = e_1^-(W')$, and the pass through v containing e has not been marked (even if $W(v) \neq W'$, as v is an articulation point in that case). The only case left, by Condition 1, is when x is the pseudo-mouth of W', because by Condition 2, x is neither the mouth nor the pseudo-mouth of $W = W(x)$; then, $x = x_k(W')$ where $k = |P(W')|$, $v = x_{k-1}(W')$ and $e = e_k^-(W')$. But $e_k^-(W'), e_{k-1}^-(W')$ where represented in $G^{\vartriangleleft}$, so they lie in the same pass through v, which, according to our marking rules, was not 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 $\tau : S^1 \to |G^2|$. 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 8 we easily have $|G^2| \approxeq |G|$, and as $|G| \approxeq |G^\sharp|$ and, trivially, $|G^\sharp| \approxeq |G^{\between\between}|$, it follows that τ is continuous and injective at ends. 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 10, which implies Theorem 3.

A finite graph G is *Hamilton-connected*, if for every two vertices x, w there is an x-w–path containing all the vertices of G. Říha [32] proved that the square of a 2-connected finite graph G is Hamilton-connected, and this fact also generalises to locally finite graphs. Indeed, if x, w are vertices of a 2-connected locally finite graph G, then adding a new vertex y^* to G, joining it to x and w by edges, and applying Theorem 10 yields a Hamilton circle from which we can delete y^* and its incident edges to obtain an x-w–arc in $|G|$ containing all the vertices of G. It is natural to ask if this remains true if we allow x, w to be ends of G, and indeed it does:

Corollary 11. *Let G be a 2-connected locally finite graph, and let $x, w \in V \cup \Omega$; then there is in $|G^2|$ an x-w–arc containing all the vertices and ends of G.*

Corollary 11 can be proved by modifying the proof of Theorem 3, so rather than proving it formally I will only point out the required modifications.

We proved Corollary 11 for the case that $x, w \in V$ above, so we may assume that w is an end. If x is a vertex then choose $y^* = x$. No further changes need to be made to the construction of the scaffolding $G^\sharp$, but instead of making it eulerian we have to give it the property that a finite cut is odd if and only if it separates x from w. This can be achieved by first making all cuts even, as we did in Section 7.3, and then choosing a ray or double ray R from x to w that contains no anchor, and replacing all single edges of R with double edges and vice-versa to obtain the new $G^\between$. The procedure of splitting into larvae and cleaning-up articulation points remains unchanged, but instead of proving that G^y is eulerian, which we did at the end of Section 7.5, we have to show that a finite cut in G^y is odd if and only if it separates x from w. In order to do so, add a new copy of $E(R)$ to both $G^\between$ and G^y to obtain auxiliary multigraphs $\dot{G}^\between$ and $\dot{G}^y$. Then note that $\dot{G}^\between$ is eulerian, and imitate the proof at the end of Section 7.5 to prove that $\dot{G}^y$ is also eulerian, from which it follows that G^y has the required property. Finally, instead of applying Theorem 4 we have to apply Corollary 8, but the rest of the proof remains unchanged.

8 Non-Locally-Finite Graphs

In this paper we proved that a locally finite graph has a Hamilton circle if it is the square of a 2-connected or the cube of a connected graph. We can also ask if this remains true for countable non-locally-finite graphs, using our definitions of Section 2 also for such graphs:

Conjecture 1. *If G is a countable 2-connected graph then $|G^2|$ contains a Hamilton circle.*

For the case that G^2 is 1-ended, Conjecture 1 has already been posed by Nash-Williams [27].

Conjecture 2. *If G is a countable connected graph then $|G^3|$ contains a Hamilton circle.*

A necessary condition for a graph G, finite or infinite, to have a Hamilton circle is that G be 1-tough — a graph G is *k-tough* if for any finite non-empty set S of vertices of G, the number of components of $G - S$ is at most $\frac{|S|}{k}$. It is easy to check that the graphs in Conjectures 1 and 2 do fulfill this condition: if G is connected then G^3 is 1-tough, and if G is 2-connected then G^2 is even 2-tough.

Let me remark that, in contrast to a locally finite graph, if G is non-locally-finite then the end-spaces of different powers of G may differ. For example, if T is the ω-regular tree, then $\Omega(T^2)$ contains the ends of T as well as a set of new ends, one for each vertex of T, but $\Omega(T^3)$ consists of one end only. Thus, although $G^k \subseteq G^{k+1}$ holds for every k, it is not clear whether G^{k+1} must be hamiltonian if G^k is.

9 Infinite Cayley Graphs

As mentioned in the introduction, it is a well-known conjecture that every finite connected Cayley graph has a Hamilton cycle. In view of Corollary 6 it is thus natural to ask if every connected locally finite Cayley graph has a Hamilton circle, however regular trees are easy counterexamples. As mentioned in Section 8, a necessary condition for a graph G, finite or infinite, to have a Hamilton circle is that G be 1-tough. Thus, an easy way to obtain infinite Cayley graphs with no Hamilton circle is by amalgamating more than k groups over a subgroup of order k. It would be interesting to decide if all non-hamiltonian connected locally finite Cayley graphs can be obtained this way:

Problem 2. *Let G be a connected Cayley graph of a finitely generated group Γ. Prove that G has a Hamilton circle unless there is a $k \in \mathbb{N}$ such that Γ is the amalgamated product of more than k groups over a subgroup of order k.*

The following problem could be used as a first step towards the solution of the previous one:

Problem 3. *Does every connected 1-ended locally finite Cayley graph have a Hamilton circle?*

10 Final Remarks

We saw that Fleischner's Theorem holds for 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 [7, 9], where Tutte's Theorem [31], 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 4 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 a suitable Euler tour in some auxiliary graph, and then try to generalise the proof to the infinite case using Theorem 4. Some open problems where this approach could be pursued are given in [22].

The following easy corollary of Theorem 4 is perhaps an argument in favour of this approach:

Corollary 12. *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

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

is well defined, and it is a bijection.

Now let σ be an Euler tour of G, that is injective at ends and maps a closed interval on each vertex of G. Let $\sigma' : S^1 \to |L(G)|$ 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 $\tau : S^1 \to |L(G)|$. 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. $\square$

Acknowledgement

I would like to thank Matthias Kriesell, and especially Philipp Sprüssel, for carefully reading an earlier draft of this document and suggesting several improvements.

References

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

[2] L. Babai. Automorphism groups, isomorphism, reconstruction. In R.L. Graham, M. Groetschel, and L. Lovász, editors, *Handbook of Combinatorics.* Elsevier, 1996.

[3] H. Bruhn. The cycle space of a 3-connected locally finite graph is generated by its finite and infinite peripheral circuits. *J. Combin. Theory (Series B)*, 92:235–256, 2004.

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

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

[6] H. Bruhn and M. Stein. On end degrees and infinite circuits in locally finite graphs. *Combinatorica*, 27:269–291, 2007.

[7] H. Bruhn and X. Yu. Hamilton circles in planar locally finite graphs. To appear in *SIAM. J. Discrete Math*.

[8] S.N. Burris. *Logic for Mathematics and Computer Science*. Prentice Hall, 1998.

[9] Q. Cui, J. Wang, and X. Yu. Hamilton circles in infinite planar graphs. To appear in *J. Combin. Theory (Series B)*.

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

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

[12] R. Diestel. Extremal infinite graph theory? a topological approach. *Oberwolfach reports*, 3:19–21, 2006.

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

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

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

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

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

[18] H. Fleischner. The square of every two-connected graph is hamiltonian. *J. Combin. Theory (Series B)*, 16:29–34, 1974.

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

[20] A. Georgakopoulos. A short proof of Fleischner's theorem. Preprint 2006.

[21] A. Georgakopoulos. Connected but not path-connected subspaces of infinite graphs. *Combinatorica*, 27(6):683–698, 2007.

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

[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*, 40:111–114, 1974.

[25] P. Heinrich. Abzählbar-unendliche Graphen mit HAMILTONscher dritter Potenz. *Math. Nachr.*, 82:231–239, 1978.

[26] J. Karaganis. On the cube of a graph. *Canad. Math. Bull.*, 11:295–296, 1968.

[27] C.St.J.A. Nash-Williams. Unexplored and semi-explored territories in graph theory. In F. Harary, editor, *New Directions in the Theory of Graphs*, pages 149–186. Academic Press, New York, 1973.

[28] E. Rapaport-Strasser. Cayley color groups and hamilton lines. *Scripta Math.*, 24:51–58, 1959.

[29] M. Sekanina. On an ordering of the set of vertices of a connected graph. *Publ. Fac. Sc. Univ. Brno*, 412:137–142, 1960.

[30] C. Thomassen. Hamiltonian paths in squares of infinite locally finite blocks. *Annals of Discrete Mathematics*, 3:269–277, 1978.

[31] W.T. Tutte. A theorem on planar graphs. *Trans. Am. Math. Soc.*, 82:99–116, 1956.

[32] S. Říha. A new proof of the theorem by Fleischner. *J. Combin. Theory (Series B)*, 52:117–123, 1991.

Agelos Georgakopoulos `http://www.math.uni-hamburg.de/home/georgakopoulos/`

Mathematisches Seminar
Universität Hamburg
Bundesstraße 55
20146 Hamburg
Germany