HAMBURGER BEITRÄGE
ZUR MATHEMATIK

Heft 249

Infinite Hamilton Cycles in Squares of Locally Finite Graphs

Angelos Georgakopoulos, Hamburg

Juni 2006

Infinite Hamilton Cycles in Squares of Locally Finite Graphs

Angelos 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 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. A well known theorem of Fleischner states that

Theorem 1 (Fleischner [7]). *If G is a finite 2-connected graph, then G^2 has a Hamilton cycle.*

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

Theorem 2 (Thomassen [15]). *If G is a 2-connected locally finite 1-ended graph, then G^2 contains both a spanning ray and a spanning double ray.*

In a 1-ended graph, it makes sense to think of a double ray as an infinite cycle: imagine that the two rays of the double ray "meet" at the only end of the graph. But what if G has more than one end? Diestel and Kühn (see [5]) introduced a natural notion of an infinite cycle in arbitrary locally finite graphs, called a *circle*: a homeomorphic copy of the unit circle S^1 in $|G|$, the topological space consisting of G and its ends (see Section 2 for details). A *Hamilton circle*, then, is a circle that traverses every vertex exactly once.

Settling a conjecture of Diestel [4, 5], we will extend Fleischner's Theorem (and Thomassen's) to locally finite graphs with any number of ends:

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

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

Theorem 4. *If a locally finite multigraph has an Euler tour, then it also has an end-faithful Euler tour.*

As discussed in Section 8, Theorem 4 could help generalise other sufficient conditions for the existence of a Hamilton cycle.

We also fully generalise to locally finite graphs the well known fact that the third power of any connected finite graph is hamiltonian:

Theorem 5. *If G is a connected locally finite graph, then G^3 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 [10].

This paper is structured as follows. After the definitions (Section 2) and some basic results (Section 3) are given, Theorem 5 is proved in Section 4. Reading this proof can be a good warming-up exercise before going into the much more complicated proof of Theorem 3, sketched in Section 5 and completed in Section 7. Theorem 4 and a further fundamental fact are proved in Section 6. Finally, Section 8, contains a corollary, some concluding remarks and a conjecture.

2 Definitions

Unless otherwise stated, we will be using the terminology of [5] for graph theoretical concepts, that of [1] for topological concepts, and that of [3] for logical ones. Let $G = (V, E)$ be a locally finite — i.e. every vertex has a finite degree — multigraph 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 xu, xv with a u-v–edge; the new edge *shortcuts* the edges xu, xv.

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

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

Apart from its usual meaning, we will, with a slight abuse, also use the term *edge* to denote the set of parallel edges between to fixed vertices in a multigraph. A *muldigraph* is a multigraph containing no triple of parallel edges. A *double edge* in a muldigraph is a pair of parallel edges; a *single edge* is an edge having no parallel.

An *arc analysis* of a finite H–bridge B, is a subgraph of B spanning $V(B - H)$, that consists of a sequence $C_1, C_2, \ldots, C_n$ of paths called *arcs*, 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 $i > 1$
- for every i, C_i contains a vertex $y(C_i) \neq p_i, q_i$ all of whose neighbours lie in $H \cup \bigcup_{j \leq i} C_j$

The endedges of C_i are its *bonds*, and $y(C_i)$ is its *articulation point.*

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

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

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, S in G are *equivalent* if no finite set of vertices separates them; we denote this fact by $R \approx_G S$, or simply by $R \approx S$ 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 $C = C(S, \omega)$ denote the component of $G - S$ that contains some (and hence a tail of every) ray in ω, and let $\Omega(S, \omega)$ denote the set of all ends of G with a ray in $C(S, \omega)$. As our basis of open neighbourhoods of ω we now take all sets of the form

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

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

It can be proved (see [6]) that in fact $|G|$ is the Freudenthal compactification [8] 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).

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.

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

[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.

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 [5, 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, but their generalisation to multigraphs is trivial).

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 is a consequence of [4, Theorem 5.4] and Lemma 2.

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

- *G is eulerian;*
- *$E(G) \in \mathcal{C}(G)$;*
- *$|E(G) \cap F|$ is even for every finite cut F of G.*

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 [9]. 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 [5, 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.*

We will make use of the compactness theorem for propositional logic (see [3]):

Theorem 6. *Let K be an infinite set of propositional formulas, every finite subset of which is satisfiable. Then K is satisfiable.*

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 matters are more complicated. In what follows, we develop some tools that will help us work with the ends of H as if they were the ends of G.

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

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

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, S in H, if $R \approx_G S$ then $R \approx_H S$. Then π_{HG} is a homeomorphism between $\Omega(H)$ and $\Omega(G)$.*

Proof. Clearly, π_{HG} is injective. Let us show that it is surjective. For any $\omega \in \Omega(G)$, pick a ray $R \in \omega$. Since H is connected, we can apply Lemma 5 to obtain a comb in H with teeth in $V(R)$. The spine of this comb is a ray in H, that is equivalent to R in G . Thus its end is mapped to ω by π_{HG}.

Since $H \subseteq G$, it follows easily that π_{HG} is continuous. Moreover, $\Omega(H)$ is compact, because it is closed in $|H|$ and $|H|$ is compact (see [5, 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, S in H, if $R \approx_G S$ then $R \approx_H S$. Let $(v_i)_{i \in \mathbb{N}}$ be a sequence of vertices of $V(G)$. Then v_i converges to $\omega \in \Omega(H)$ in $|H|$ if and only if v_i converges to $\pi_{HG}(\omega)$ in $|G|$.*

Proof. Define a mapping $\hat{\pi}_{HG} : V(H) \cup \Omega(H) \to V(G) \cup \Omega(G)$ 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 $H \subseteq G$, 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, S in H, such that $R \approx_G S$. By Lemmas 6 and 7, it suffices to show that $R \approx_H S$.

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

4 The Third Power of a Locally Finite Graph is Hamiltonian

It is a well known fact, that the third power of a connected finite graph is hamiltonian. Extentions of this fact to infinite graphs have been made by Sekanina

[14], who showed that the third power of a connected, locally finite, 1-ended graph has a spanning ray, and by Heinrich [12], 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 [5, Theorem 8.2.4]), and let $T_i = T[v]^i$.

We will prove the assertion using Theorem 6. To this end, define for each edge $e \in E(T^3)$ a logical variable $v(e)$, the truth-values of which encode the presence or not of e, and let $\mathcal{V}$ be the set of these variables. For every vertex $x \in T$, write a propositional formula with variables in $\mathcal{V}$, expressing the fact that exactly two x-edges are present, and let $\mathcal{P}_1$ be the set of these formulas. For every branch B of T, write a propositional formula with variables in $\mathcal{V}$ expressing the fact that at most two edges that cross B are present, and let $\mathcal{P}_2$ be the set of these formulas. For every finite cut F of T^3, write a propositional formula with variables in $\mathcal{V}$, expressing the fact that an even, positive number of edges in F are present, and let $\mathcal{P}_3$ be the set of these formulas. Let $\mathcal{P} = \mathcal{P}_1 \cup \mathcal{P}_2 \cup \mathcal{P}_3$.

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

Lemma 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 that two edges of H, if this is true for the Hamilton cycles H_v of the v–branches. □

So by Theorem 6, 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, $F \in \mathcal{C}(T^3)$, 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 $V(C) \nsubseteq V(T_i)$, let $\mathcal{B}$ be the (non-empty) set of branches B in $T - T_i$ such that $B \cap V(C) \neq \emptyset$, and let $X = V(C) \cup \bigcup_{B \in \mathcal{B}} V(B)$. Since $u \notin X$, $E(X, X' := V(T) - X)$ is a non-empty cut D, which is clearly finite. Now for every $B \in \mathcal{B}$, there is a formula in $\mathcal{P}_2$ asserting that there are at most two edges crossing B, and since (by Lemma 4 and Lemma 2) C already contains two such

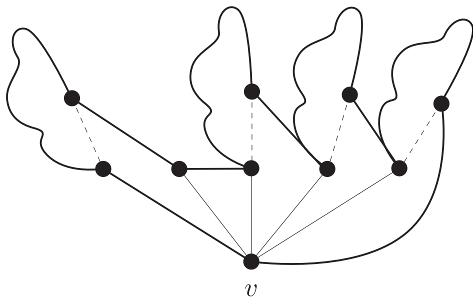

Figure 1: Using the induction hypothesis to pick a Hamilton cycle of T. The wavy curves represent Hamilton cycles of the v-branches supplied by the induction hypothesis, and for each such Hamilton cycle H, $e(H)$ is represented by a dashed line. The thick cycle represents H.

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 $|G^3| \approxeq |G|$. Easily by Lemma 7, $|T| \approxeq |G|$, thus H is also a Hamilton circle of G^3. $\square$

5 Outline of the Proof of Theorem 3

Before giving an outline of the proof of Theorem 3, let me compare it with the proofs of Theorem 7, which is a strengthened form of Theorem 1 introduced in Section 7.1, and Theorem 2. The descriptions that follow are approximate, omitting much information not needed for the comparison. Riha [13] proves Theorem 7 by induction; he finds a special cycle $C \ni y^*$, and then applies the induction hypothesis to each component of $G - C$, to obtain a set of C–paths in G^2, called *basic paths*, so that each vertex of $G - C$ lies in exactly one of the paths. Basic paths have the property that their endedges are original edges of G; let us call these endedges bonds. This property makes it possible, to recursively merge pairs of incident basic paths into longer basic paths, by shortcutting incident bonds, and he repeats this operation as often as possible without disconnecting the graph. Then, some edges of C are replaced by double edges, so that the resulting multigraph is eulerian. Finally, it is shown that every Euler tour J of this multigraph, can be transformed to a Hamilton circle of G^2, by replacing some subtrails of length two of J, with edges of G^2 with the same endvertices; we call this process *hamiltonisation* of J.

Thomassen follows a similar plan in his proof of Theorem 2 (which appeared before Riha's proof). The cycle C is replaced by a ray R, such that all components of $G - R$ are finite, and Theorem 7 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 Riha'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.

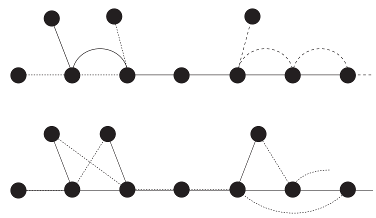

Figure 2: An example of a local hamiltonisation. In the upper figure, the restriction of the Euler tour on the segment (horizontal path) is indicated; it consists of three paths. In the figure beneath, these trails have been transformed to disjoint paths in the square of the graph, that span all vertices (dashed lines).

Trying to imitate these proofs for arbitrary locally finite graphs, we face three major problems. The first one regards Euler tours. In the sketched proofs, an Euler tour was transformed to a Hamilton circle, by performing "leaps" over one vertex, using an edge of G^2. Doing so for an arbitrary Euler tour of a locally finite graph, we cannot avoid running through some end more than once. But a Hamilton circle must, by definition, traverse each end exactly once, thus if we want to gain one from an Euler tour using this method, the Euler tour itself should be end-faithful. So we have to ask, which eulerian graphs admit an end-faithful Euler tour. The answer is given by Theorem 4: all of them.

The second problem is, what the analogue of R or C should be. Deleting any ray will, in general, leave infinite components behind, and we cannot easily have an induction hypothesis to replace them by basic paths. Instead, we will use a complicated structure, looking like an end-faithful 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 perform all possible shortcuts in this graph, we will end up with a 2-ended graph G', because the basic paths will merge into a ray non-equivalent with R. We can still continue with the plan of finding an Euler tour and transforming it to a Hamilton circle H of G', but H will not be a Hamilton circle of G: it will traverse ω twice.

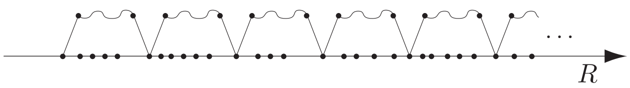

Figure 3: Performing all shortcuts between bonds would create a new end.

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 can 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 by 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. So on the one hand we should avoid shortcuts because they are dangerous for the end topology, and on the other we need them in order to get rid of vertices like x. An equilibrium is needed, which I could not find.

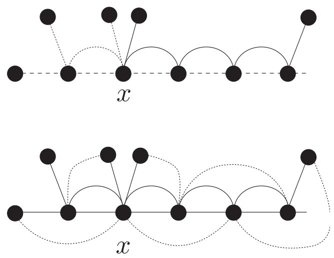

Figure 4: A difficult case: the lower figure shows the only possible hamiltonisation of the trails shown in the upper figure.

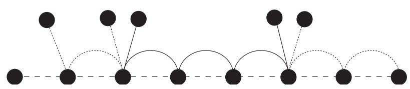

Figure 5: A case where no hamiltonisation is possible.

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 two which the Euler tour must traverse. Practically, 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 Euler tour of Figure 5 could not result from an Euler tour of the graph in Figure 6. The idea of imposing such constraints on the Euler tour is used in [10] to obtain a short proof of Theorem 1.

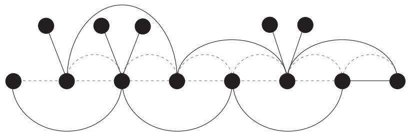

Figure 6: Applying constraints for an Euler tour of the graph of Figure 5. The dashed lines represent edges of the original graph not in the auxiliary graph.

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, called rope-ladders, like the one shown in Figure 7, irregularly attached on each other. The finite components that the rope-ladders leave behind, are covered by arc analyses.

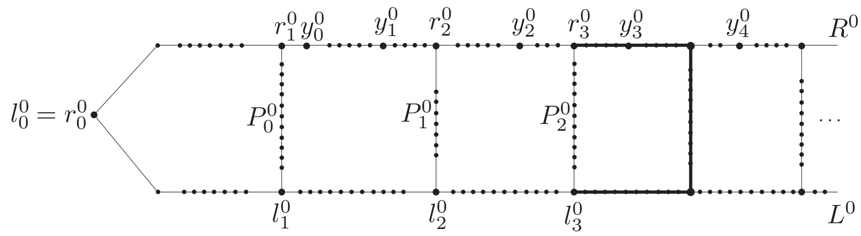

Figure 7: A rope-ladder. The horizontal paths are equivalent rays; usually their first vertices do not coincide.

Making the graph eulerian 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 the graph 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 consider three cases. If one of the endedges of P is a double edge — we know that at most one is — we delete it, and consider the rest of P as a larva. If not, then we look at a special vertex in P, denoted by y_j^i, and called an articulation point, and if one of the edges of P incident with the articulation point is double, we delete it, and consider the two remaining subpaths of P as larvae. If both edges incident with the articulation point are single, we consider the two maximal subpaths of P ending at the articulation point as larvae (Figure 8). The first pi of each rope-ladder however, does not follow these rules, which is the source of the anomaly regarding y_0^0 in Figure 7. An arc analysis is treated in a similar way. In all cases, we make sure that the first edge of every larva is a single edge.

Having divided the whole graph in 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 Riha and Thomassen are needed, with the exception of the articulation points. As shown in Figure 8, it could happen that two larvae intersect at their last vertex, which is an articulation point. 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

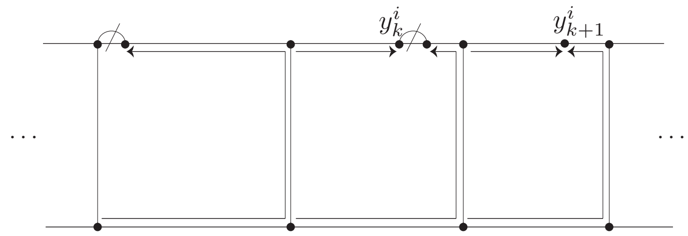

Figure 8: Splitting the graph in larvae

to shortcut this edge. If larvae intersect otherwise however, there could be a conflict between their hamiltonisations. In order to avoid this 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 Riha's proof: there, C contains a vertex with the property that it sends no edge to the rest of the graph, and this vertex has 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, where we will, in Section 7.6, pick the end-faithful Euler tour. Of course, we will need to prove that the auxiliary graph bears the same end topology as the original one. Then, based on the fact that the Euler tour complies with the constraints we imposed on it, we will show that it is possible to hamiltonise it in Section 7.6.

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

1. construct the scaffolding;
2. make it eulerian;
3. split it into larvae;
4. impose constraints on the Euler tour;
5. clean up the articulation points;
6. pick an Euler tour and hamiltonise it.

6 Some Preliminaries

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 $k \in \mathbb{N}$ disjoint ω-rays, then it also has a set K' of k disjoint ω-rays that devours ω. Moreover, K' can be chosen so that its rays have the same starting vertices as the rays in K.*

Proof. We will perform induction on k. For $k = 1$ this is easy; the desired ray can for example be obtained by imitating the construction of normal spanning trees in [5, Theorem 8.2.4]. For the inductive step, let $K = \{R_0, R_1, \dots, R_{k-1}\}$ be a set of disjoint ω-rays in G. We want to apply the induction hypothesis to $G - R_0$, but we have to bear in mind that some of the R_i might not be equivalent to one another after deleting R_0. So let $S \subset V$ be a finite set, so that any two tails of elements of K that lie in the same component of $G - R_0 - S$ are equivalent. Applying the induction hypothesis to all components of $G - R_0 - S$ that contain a tail of some R_i, we obtain a new set of rays $R'_1, R'_2, \dots, R'_{k-1}$ so that any ray equivalent to some R_i in $G - R_0 - S$ meets some R'_j, and for each j, R'_j starts at the first vertex of R_j not in S. We can now prolong each R'_j using the subpath of R_j that lies in S, to achieve that R'_j and R_j start at the same vertex (without loss of generality, each R_j leaves S only once, because otherwise we can add an initial subpath of R_j to S). Moreover, let R'_0 be a ray in $G - \{R'_1, R'_2, \dots, R'_{k-1}\}$ meeting all rays equivalent with R_0 in that graph and starting at the first vertex of R_0. We claim that $K' = \{R'_0, R'_1, \dots, R'_{k-1}\}$ meets every ω-ray in G.

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

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

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

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 be 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 an end-faithful topological Euler tour σ of G. After each step i, we will have defined a finite set of edges F_i, and a continuous and onto function $\sigma_i : S^1 \to \bar{F}_i$, 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 from $E(G) - F_i$, put them in F_i to obtain F_{i+1}, and modify σ_i to σ_{i+1}. We might also add some vertices to S_i to obtain S_{i+1}.

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

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

So we may define

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

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 $F_i, i \in \mathbb{N}$ is converging, so it suffices to consider the basic open sets of that form, and it is easy to see that their preimages are indeed open.

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

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

7 Proof of Theorem 3

7.1 A Stronger Assertion

A shorter proof of Theorem 1 was given by Riha [13], who in fact proved a slightly stronger assertion:

Theorem 7. *If G is a finite 2-connected graph, $y^* \in V(G)$ and $e^* = y^*x^* \in E(G)$, then G^2 has a Hamilton cycle H that contains e^* and a y^*-edge $e' \in E(G), e' \neq e^*$.*

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

Theorem 8. *If G is an infinite 2-connected locally finite graph, $y^* \in V(G)$ and $e^* = y^*x^* \in E(G)$, then G^2 has a Hamilton circle H that contains e^* and a y^*-edge $e' \in E(G), e' \neq e^*$.*

7.2 Constructing the Scaffolding

In this section we construct the "scaffolding" $G^\sharp$ mentioned in Section 5. The following lemma is similar to a lemma of Riha [13].

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 arc analysis such that x lies in C_1.*

Proof. Pick an H-path C starting at x, and let D be a component of $B-(C\cup H)$; if there is no such component, then we can let $C_1 = C$, pick any inner vertex of C_1 as $y(C_1)$, and choose C_1 as an arc analysis of B. Suppose that C, D have been chosen so that $|V(D)|$ is minimal. Clearly, D has at least one neighbour u on $C - H$. If it has more than one, then let P be a subpath of $C - H$ whose endvertices u, v are neighbours of D, such that no inner vertex of P is a neighbour of D, and let C_1 be the path resulting from C after replacing P with a v-u–path with all inner vertices in D. If u is the only neighbour of D on $C - H$, then let v be a neighbour of D in H, and replace one of the subpaths of C connecting u to H, with a v-u–path with all inner vertices in D (having at least one inner vertex), so that the resulting path C_1 meets x.

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

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

By a result of Halin [11, 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 $(x_i)_{i\in\mathbb{N}}$ be an enumeration of $V := V(G)$. Let ω be any end of G. By Lemma 13, there are two disjoint ω-rays starting at x^* and y^*, and by Lemma 10 there are rays R^0, L starting at y^*, x^* respectively, that devour ω. Let $L^0 = y^*x^*L$, and let $r^0_0, l^0_0 = y^*$. Choose a sequence $(y^0_j)_{j\in\mathbb{N}}$ 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 on $\mathring{r}^0_{j-1} R^0 \mathring{r}^0_{j+1} \cup \mathring{l}^0_{j-1} L^0 \mathring{l}^0_{j+1} - y^0_0$

All these conditions are easy to satisfy, if we choose the y^0_j and P^0_j in the order $P^0_0, y^0_1, P^0_1, y^0_2, P^0_2, \ldots$: 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. Note that by Lemma 11 every $(R^0 \cup L^0)$–bridge has finitely many feet.

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 an arc analysis of every finite RL^0_-–bridge, which exists by Lemma 12. Let $G^\sharp_0 = RL^0$.

The construction of $G^\sharp_0$ was the first step in an induction whose aim is to define $G^\sharp$. Each step i of this induction 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 finite bridges to $G^\sharp_{i-1}$ to obtain $G^\sharp_i$. The endpoints of R^i, L^i will be distinct vertices of $G^\sharp_{i-1}$. The structure of $G^\sharp$ will resemble that of an end-faithful spanning tree of G, where instead of rays we have structures like RL^0.

Formally, 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-RL^l–path in $\hat{C}_i$, where l is the greatest index for which such a path exists. Let $v = v(i)$ be the last vertex of Q_j not in $G^\sharp_{i-1}$, and $w = w(i)$ the vertex after v on Q_j (thus $w \in G^\sharp_{i-1}$). We claim that:

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

Proof. Contracting $G - C_i$ to one vertex z, we obtain a 2-connected graph, in which we can apply Lemma 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 on $\mathring{r}^i_{j-1} R^i \mathring{r}^i_{j+1} \cup \mathring{l}^i_{j-1} L^i \mathring{l}^i_{j+1} - y^i_0$

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 an arc analysis of every finite $RL^i\_$–bridge in G. We call RL^i a *rope-ladder* (RL^0 is also a rope-ladder). 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 $y^* x^*$ 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 $r^i_0 R^i P^i_0 L^i l^i_0$, 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, its endedges). Call the edges of RL^0_- incident with y^* the *bonds* of RL^0. Recall that arcs also have bonds and articulation points. The following assertion is true by construction:

Observation 1. *If RL^i sends a bond to $RL^j - r^j_0 - l^j_0$, then $RL^j \preceq RL^i$*

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 arc of an arc analysis, 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 arcs (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 arc and it sends a bond to an inner vertex of Π (y^* is an inner vertex of Π^0_0). Let $\preceq$ be the reflexive transitive closure of $\prec$. Clearly, $\preceq$ is a partial order.

Define $G^\sharp_i$ as the union of $G^\sharp_{i-1}$ with RL^i and an arc analysis 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 vary from G too much.

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

Our next aim is to prove that $|G^\sharp| \cong |G|$. Suppose Q, T are rays in $G^\sharp$ so 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. Thus there is a j so that $G^\sharp_j$ separates Q from 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 let $q \in V(Q) \cap C \cap RL^k_-$ 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 $R \preceq RL^k, RL^l$. But then, we can find a t-R–path P_1 and a q-R–path P_2 in $G^\sharp$, that use only vertices of rope-ladders RL^i such that $R \preceq RL^i \preceq RL^l$ and $R \preceq RL^i \preceq RL^k$ respectively. The union of P_1, P_2 lies 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, $|G| \cong |G^\sharp|$.

7.3 Making the Graph Eulerian

The next step is to change $G^\sharp$ to an Eulerian muldigraph $G^\between$, where all anchors are single edges. Rather than constructing the muldigraph explicitly, we will show its existence using Theorem 6. In order to meet its condition, we will show that:

Claim. *For every $i \in \mathbb{N}$ there is an Eulerian muldigraph $G_i^\between$ on V, so that x, y are neighbours in $G_i^\between$ if and only if they are neighbours in $G^\sharp$, and furthermore every anchor in $G^\sharp[y^*]^i$ is a single edge in $G_i^\between$.*

Proof. If C is a cycle in the muldigraph G, then *switching* C is the operation of replacing each double edge of C with a single edge, and each single edge of C with a double edge.

In order to prove the Claim, begin by doubling all edges of $G^\sharp$. Then, for every arc analysis $C_1, C_2, \dots, C_k$ meeting $G^\sharp[y^*]^i$, recursively, for $j = k, k-1, \dots, 0$, if the anchor of C_j is a double edge, find a cycle containing C_j and avoiding $\bigcup_{i>j} C_i$ and all other arc analyses in $G^\sharp$, and switch this cycle. After doing so for all arc analyses, recursively for $j = l, l-1, \dots, 1$, where l is the greatest index such that the anchor of RL^l lies in $G^\sharp[y^*]^i$, if the anchor of RL^j is a double edge, switch a cycle comprising Π_0^j and a path in $G_{j-1}^\sharp$ that has the same endvertices as Π_0^j and contains no edge of an arc analysis (for $j = 0$ switch Π_0^0). After you are done with 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 muldigraph. Note that $G_i^\between$ resulted from a muldigraph where all edges 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. Obviously, $G_i^\between$ satisfies all conditions of the Claim. $\square$

In order to apply Theorem 6, 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 single.

By Theorem 6 and the Claim, there is an assignment of truth-values to the elements of $\mathcal{V}$ satisfying all these propositional formulas. This assignment encodes an assignment of multiplicities to the edges of $G^\sharp$, which defines a muldigraph $G^\between$ which is Eulerian (by Lemma 3), and in which all anchors of $G^\sharp$ are single edges.

Let $G^{\not\between}$ be the muldigraph resulting from $G^\between$ after deleting all bonds that are double edges. We claim that $|G^{\not\between}| \approxeq |G^\sharp|$. 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 $j = \min(\{k < i | e \cap RL^k \neq \emptyset\})$. Note that $RL^l \cap G^{\not\between}$ is connected for every l, and if RL^i is anchored on RL^k, then for any two vertices $r \in RL^i$, $t \in RL^k$, there is an r-t–path P in $(RL^i \cup RL^k) \cap G^{\not\between}$, and if both bonds of RL^i meet RL^k, then P can be chosen so that it avoids every window of RL^k that has a lower index than the lowest window of RL^k met by a bond of RL^i. By Observation 1, $RL^j \preceq RL^i$, so we can use the latter fact recursively to obtain a p-q–path $d(e)$ in $G^{\not\between}$ that avoids $G^\sharp_{j-1}$ and the windows of RL^j up to the first one that meets RL^i. If e is a deleted bond of an arc of an arc analysis D, we can, by a similar argument, find a p-q–path $d(e)$ in $G^{\not\between}$ avoiding $G^\sharp_{j-1}$ and the windows of RL^j up to the first one that meets D, where j is the least index of a rope-ladder meeting D. Finally, if e is a deleted bond of a window W, then let $d(e) = W - e$.

The set $\{d(e)|e \in E(G^{\between}) - E(G^{\not\between})\}$ is thin. To see this, note that for any edge $f \in \Pi^i_j$ for some i, j, only those rope-ladders and arc analyses sending a bond to some RL^k with $k < i$, or to some Π^i_k with $k \leq j$ can contribute a $d(e)$ containing f, and there are only finitely many such rope-ladders and arc analyses. Thus by Lemma 8, $|G^{\not\between}| \approxeq |G^{\between}| \approxeq |G^\sharp|$.

7.4 Splitting into Finite Multigraphs

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 in $G^{\not\between}$, s is one of its endvertices, called its *mouth*, and the edge of P incident with s is single. For every larva $W = (s, P)$, we label the vertices of P with $x_i = x_i(W)$, so that $P = x_0(= s)x_1x_2 \ldots x_n$. Moreover, let $e_i = e_i(W)$ denote the edge $x_{i-1}x_i$, and if e_i is a double edge denote its copies by $e_i^-, e_i^\cap$, otherwise let $e_i^- = e_i$. Let $P(W) = P$. Whenever we use an expression assuming a direction on P or W, we consider x_0 to be its first vertex and x_n its last. In order to simplify the notation, we will also write sPy instead of (s, sPy).

Recall that we want to impose some constraints on the Euler tour that is supposed to produce a Hamilton circle of G. This is done separately for each larva following the pattern of Figure 6: *metamorphosing* the larva $W =: (s, P)$, is the operation of replacing, in P and in $G^{\not\between}$, the edges $e_{i+1}^\cap, e_i^-$, for every i such that e_{i+1} is a double edge, by an $x_{i-1}x_{i+1}$ edge f_i, called a *representing* edge. Note that $e_{i+1}^\cap, e_i^-, f_i$ form a triangle. The *caterpillar of* W is the graph X resulting from P after metamorphosing W. Note that X is connected. Each time we metamorphose a larva, we will assume that for each edge e deleted, a detour $dt(e)$ for e is specified in X.

If the last edge e_k of P is a single edge, then *completely metamorphosing* W is the operation of metamorphosing W, and then replacing e_k^-, e_{k-1}^- with an $x_{k-2}x_k$ edge f_{k-1}, also called a *representing* edge. If W is completely metamorphosed, then its *pseudo-mouth* is its last vertex. The *double caterpillar of* W is the graph X resulting from P after completely metamorphosing W. A double caterpillar has a big advantage in comparison to a caterpillar: the additional constraint (on the Euler tour), allows it to be hamiltonised so that its last edge, as well as its first, is not shortcutted, and so its pseudo-mouth is allowed to meet other larvae (even if it is not an articulation point). This advantage however, comes at a high price: a double caterpillar is a disconnected graph, with two components. For this reason, each time we completely metamorphose a larva W to obtain X, we will specify some *detour* $dt(X)$ for X, that is, an X-path connecting the two components of X (note that the last two vertices of $P(W)$ lie in distinct components of X, and in fact $dt(X)$ will always be a path connecting those vertices). We assume that also for each edge e deleted while completely metamorphosing W to get X, a detour $dt(e)$ for e is specified in $X \cup dt(X)$.

We now divide the graph in larvae, and either metamorphose or completely metamorphose each of them. More precisely, we will specify a set of edge-disjoint larvae $\mathcal{W}$ so that $G^{\not\propto} = \bigcup_{W \in \mathcal{W}} P(W)$, and the following conditions are satisfied:

Condition 1. *If $W, W' \in \mathcal{W}$, then W, W' are edge-disjoint, and if $v \in P(W) \cap P(W')$ then one of the following is the case:*

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

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

We will construct a muldigraph $G^{\lhd}$ on V by performing operations of the following kinds on $G^{\not\propto}$:

- replacing two edges e, f with an edge forming a triangle with e, f;
- switching a window;
- adding a double edge from $G^{\propto} - G^{\not\propto}$;
- 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 $G^{\lhd}$, 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 arc 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 spanned by RL^i_- connected for every i, so as to guarantee that the end topology remains the same. If however, we apply those rules to Π_0^i, then we could disconnect part of it from the rest of RL^i_-. To avoid this, we will treat pis of the form Π_0^i differently.

So we will construct $G^{\lhd}$ 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.

For the first phase, perform ω steps of the following kind. In step i, if Π_0^i has already been handled, that is, divided in larvae, in some previous step, or if one of its bonds is not present in $G^{\not\propto}$, go to the next step. Otherwise, if a bond of Π_1^i is not present in $G^{\not\propto}$, then put it back (it is a double edge). We consider two cases.

In the first case, called Case I, both edges $e = r_1^i y, e'$ incident with $y := y(\Pi_0^i)$ on Π_1^i are single or both are double. If they are both double, then switch W_1^i. No matter if we switched W_1^i or not, metamorphose the larvae (r_1^i, e) (this is a trivial larva) and $l_0^i \Pi_0^i l_1^i \Pi_1^i y$ (see Figure 9). If the edge $d = l_1^i l'$ of P_0^i incident with l_1^i is double, delete d and metamorphose the larva $r_0^i \Pi_0^i l'$; pick a detour $dt(d)$ for d in the three resulting caterpillars. If d is single, and there is a double edge $f = uv$ on P_0^i, delete f and metamorphose the larvae $l_1^i P_0^i f$ and $r_0^i \Pi_0^i f$; pick a detour $dt(f)$ in the four resulting caterpillars. If there is no double edge on P_0^i, let r' be the neighbour of r_1^i on P_0^i, metamorphose the larva $l_1^i P_0^i r'$ and completely metamorphose the larva $r_0^i \Pi_0^i r'$ (see Figure 10); a detour for the double caterpillar can be found in the resulting caterpillars. It is easy to confirm that Condition 1 is always satisfied, and that Condition 2 is satisfied for all vertices in $\Pi_0^i - \{a(\Pi_0^i), b(\Pi_0^i)\}$. However $a(\Pi_0^i), b(\Pi_0^i)$ lie in lower pis (if $i > 0$), which are responsible for them. Moreover, the following is true:

Observation 2. *No detour specified in Case I meets any pi $\Pi \neq \Pi_0^i$ for which $\Pi \preceq \Pi_0^i$ holds.*

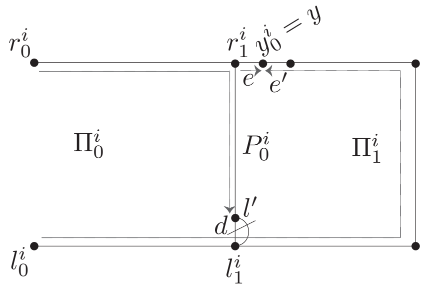

Figure 9: Splitting into larvae: Case I, and d is double. The dashed lines indicate larvae, and arrows show away from the mouth.

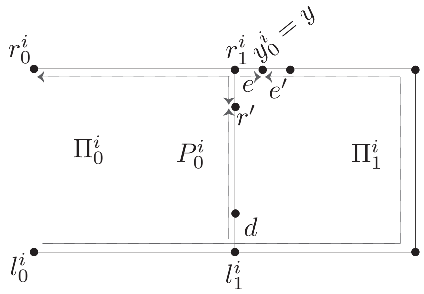

Figure 10: Splitting into larvae: Case I, and no double edge on P_0^i. The line with arrows at both ends indicates a larva that will be completely metamorphosed.

In the second case, called Case II, one of e, e' is single and the other is double. We want to choose and metamorphose some larvae, that will give rise to an RL^i_{-}–path A^i in $G^{\triangleleft}$ with one endpoint at y, which will help us delete an edge in RL^i_{-} without putting the end topology at risk.

Since y has even degree, there is at least one single bond (other than e) incident with y. Pick such a bond b, so that the pi or arc Π_0 of which b is a bond is minimal under $\preceq$. Note that the anchor of Π_0 cannot be b, since we are at an articulation point. Let Π_1 be the pi or arc 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 RL^i_-–path. If not, we go on recursively, trying in each step j to extend the already chosen initial subpath A_{j-1} of A^i, by attaching a path in $V(\Pi_j)$, where Π_j contains the endpoint of A_{j-1}, to reach a pi or arc $\Pi_{j+1} \preceq \Pi_j$. 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 $G^{\emptyset}$, 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 arc 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 Π_j has two single bonds, then there are two cases.

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

In the case that $y(\Pi_j)$ is incident with no double edge on Π_j, metamorphose the larvae $a(\Pi_j)\Pi_j y(\Pi_j)$ and $b(\Pi_j)\Pi_j y(\Pi_j)$. Let A_j be the concatenation of A_{j-1} with an A_{j-1}-$a(\Pi_j)$–path in the resulting caterpillars. Let Π_{j+1} be the pi or arc 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 RL^i_-–path with precisely one endpoint at y. We call it the *apophysis of RL^i*. If Π_{j+1} does not lie in RL^i_-, we proceed with the next step. Clearly $\Pi_{j+1} \preceq \Pi_0$, and furthermore $\Pi^i_0 \preceq \Pi_{j+1}$, 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 arcs 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 $\Pi^k_1 \prec \Pi^i_0$, it has to go through Π^i_0. 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.

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 4 is even satisfied by pis used by some apophysis):

Condition 3. *The union of A^i with the graph induced by $V(\Pi^i_0 \cup \Pi^i_1)$ after metamorphosing all larvae in $\mathcal{W}_i$ is connected.*

Condition 4. *r^i_2, l^i_2 lie in the same larva in $\mathcal{W}$.*

First we consider the case $y' \in \Pi^i_1 - r^i_1$ (Figure 11). If e is double, then switch W^i_1. Now e is single and e' double; delete e'. Then metamorphose the trivial larva (r_1^i, e), and the larva $l_0^i\Pi_0^i l_1^i\Pi_1^i e'$. Pick a detour $dt(e')$ for e' in the union of A^i with the resulting caterpillar. Handle Π_0^i like in Case I: if the edge $d = l_1^i l'$ of P_0^i incident with l_1^i is double, delete it and metamorphose the larva $r_0^i\Pi_0^i l'$; pick a detour $dt(d)$ for d in the resulting caterpillars and A^i. If d is single, and there is a double edge f on P_0^i (Figure 11), delete f and metamorphose the larva $l_1^i P_0^i f$ and the larva $r_0^i\Pi_0^i f$; pick a detour $dt(f)$ in the resulting caterpillars and A^i. If there is no double edge on P_0^i, let r' be the neighbour of r_1^i on P_0^i, metamorphose the larva $l_1^i P_0^i r'$ and completely metamorphose the larva $r_0^i\Pi_0^i r'$; a detour for the resulting double caterpillar can again be found in the resulting caterpillars and A^i.

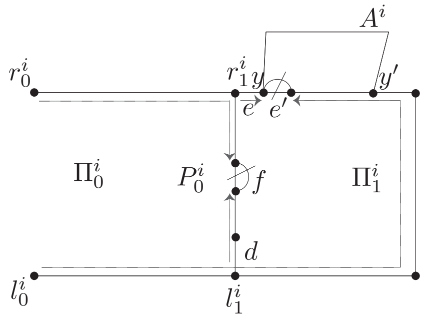

Figure 11: Splitting into larvae: Case II, $y' \in \Pi_1^i - r_1^i$, and there is a double edge f on P_0^i.

In the case that $y' \in l_0^i\Pi_0^i\mathring{l}_1^i$, switch W_1^i if needed, so that e is single and e' double; delete e'. Then metamorphose the trivial larva (r_1^i, e), and the larva $r_0^i\Pi_0^i l_1^i\Pi_1^i e'$ (Figure 12). Pick a detour $dt(e')$ for e' in the union of A^i with the resulting caterpillars. Then, if the first edge $h = l_1^i l''$ of $l_1^i\Pi_0^i y'$ is double, delete it and metamorphose the larva $l_0^i\Pi_0^i l''$; pick a detour $dt(h)$ for h in the resulting caterpillars and A^i. If h is single, and there is a double edge f on $l_1^i\Pi_0^i y'$, delete it and metamorphose the larva $l_1^i\Pi_0^i f$ and the larva $l_0^i\Pi_0^i f$; pick a detour $dt(f)$ in the resulting caterpillars and A^i. If there is no double edge on $l_1^i\Pi_0^i y'$, let z be the neighbour of y' on $l_1^i\Pi_0^i y'$, metamorphose the larva $l_1^i\Pi_0^i z$ and completely metamorphose the larva $l_0^i\Pi_0^i z$ (Figure 12); a detour for the resulting double caterpillar can again be found in the resulting caterpillars and A^i.

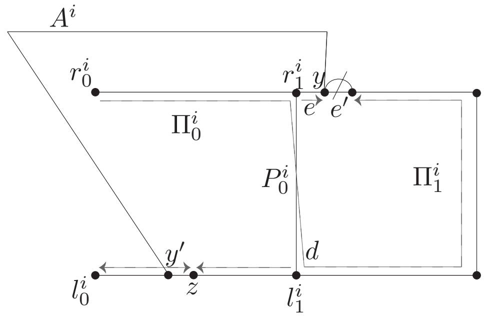

Figure 12: Splitting into larvae: Case II, $y' \in l_0^i\Pi_0^i\mathring{l}_1^i$, and no double edge on $l_1^i\Pi_0^i y'$.

Finally, if $y' \in r_0^i\Pi_0^i\mathring{l}_1^i$, switch W_1^i if needed, so that e is double and e' single;

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

In all cases, Condition 1 is satisfied. Condition 2 is satisfied for all vertices in $\Pi_0^i - \{a(\Pi_0^i), b(\Pi_0^i)\}$, but $a(\Pi_0^i), b(\Pi_0^i)$ lie in lower pis (if $i > 0$), which are responsible for them. Moreover, the following is true:

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

Now is the time to specify a detour $dt(d)$ for each edge d we deleted during the construction of A^i. It will suffice to construct paths D_1, D_2 each connecting an endpoint of d to $RL_-^i \cup A^i$. Then, since D_1, D_2 can only meet RL_-^i in Π_0^i or Π_1^i by the choice of P_0^i, P_1^i, we can, by Condition 3, find a path D in $V(\Pi_0^i \cup \Pi_1^i \cup A^i)$ connecting the endpoints of D_1, D_2, and set $dt(d) = D_1 \cup D \cup D_2$.

Deleting e separated the pi or arc 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 an e-A^i–path in the corresponding caterpillar. In order to choose D_2, we imitate the procedure we used to construct A^i: we separate the pi or arc on which Q_2 lands in one or two larvae, unless it has already been handled (that is, separated 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 nearer to RL_-^i (or A^i). We repeat until we meet $RL_-^i \cup A^i$.

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

Observation 4. *No detour for an edge deleted while constructing* A^i *meets any pi* $\Pi \neq \Pi_0^i$ *for which* $\Pi \preceq \Pi_0^i$ *holds.*

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

Observation 5. $dt(f)$ *does not meet any pi* $\Pi \neq \Pi_{l-1}^k$ *for which* $\Pi \preceq \Pi_{l-1}^k$ *holds.*

Having handled all pis, we go on to the arc analyses. For every arc analysis D with arcs $C_1, C_2, \dots, C_k$, recursively for $i = k, k-1, \dots, 1$, if C_i has not already been handled (while constructing some apophysis), then we want to split C_i in 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; precisely, we will split C_i in larvae, metamorphose them, and perhaps make some edge replacements, 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 arc $\Pi \preceq C_i, \Pi \neq C_i$, and contains no edge incident with some y^j_l.*

We consider two cases. For the first case, if $C_i \cap G^{\emptyset}$ 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 $G^{\emptyset}$, we metamorphose the larva $a(C_i)C_ib$. If not and both edges on C_i incident with $y := y(C_i)$ are single, we metamorphose the larvae $a(C_i)C_iy$ and $b(C_i)C_iy$. Otherwise, we delete a double edge f incident with y, and metamorphose the larvae $a(C_i)C_if$ and $b(C_i)C_if$. Clearly, Condition 5 is now satisfied.

In the second case, when $C_i \cap G^{\emptyset}$ meets y^j_l for some j, l, note that both bonds of C_i must be present in $G^{\emptyset}$, as by definition the anchor of C_i does not meet y^j_l. Now if both edges on C_i incident with $y := y(C_i)$ are single, metamorphose the larvae $a(C_i)C_iy$ and $b(C_i)C_iy$. Otherwise, since all vertices have even degree, there is in $\bigcup D$ some $RL^j_- \cup \bigcup_{n\leq i} C_n$–bridge B meeting $a(C_i)C_iy$ at some inner vertex u (recall that by construction, no $RL^j_- \cup \bigcup_{n\leq i} C_n$–bridge meets y). Again since all vertices have even degree, B has at least one foot $v \neq u$ in $\bigcup_{n\leq i} C_n \cup RL^j_-$; let P be a u-v–path in B. We consider four subcases:

If $v \notin C_i$, and there is a double edge f in $a(C_i)C_iu$, delete f and metamorphose the larvae $a(C_i)C_if$ and $y^j_lC_if$.

If $v \notin C_i$, and there is no double edge in $a(C_i)C_iu$, let u' be the neighbour of u on $a(C_i)C_iu$, metamorphose the larva $a(C_i)C_iu'$ and completely metamorphose the larva $y^j_lC_iu'$.

If $v \in C_i$, and there is a double edge f in vC_iu, delete f and metamorphose the larvae $a(C_i)C_if$ and $y^j_lC_if$.

If $v \in C_i$, and there is no double edge in vC_iu, suppose without loss of generality that $v \in y^j_lC_iu$. Let u' be the neighbour of u on vC_iu, metamorphose the larva $y^j_lC_iu'$ and completely metamorphose the larva $a(C_i)C_iu'$.

In the last two subcases, if in addition $v = y^j_l$ then let X be the caterpillar containing v, and shortcut the edges of P,X incident with y^j_l; call the new edge a *shortcutting edge* (note that this change does not affect the satisfaction of Condition 5 by the arcs in B, neither does it affect any apophysis; indeed, the only vulnerable apophysis is A^j, and by the choice of the Π_0 of A^j, $\Pi_0 \preceq C_i$, thus no arc in B participates in A^j). In case one of the shortcutted edges f lies on some previously specified detour d, use $dt(f)$, which will be defined below, to amend d.

In all cases, Condition 5 is now satisfied for C_i. For example, in the second subcase, $V(C_i)$ is divided in a caterpillar X and a double caterpillar Y. 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 either there is an x-u–path in Y avoiding y^j_l, which can be extended by P to an x-v–path, or there is an x-u'–path in Y avoiding y^j_l, which can be extended by a u'-$a(C_i)$–path in X to an x-$a(C_i)$–path.

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

This completes the second phase. Denote the resulting muldigraph by $G^{\triangleleft}$. Let $G_1 := (V, E(G^{\triangleleft}) \cup E(G^{\not\between}))$. Easily, by Lemma 8, $|G_1| \approxeq |G^{\not\between}|$. The set $\{dt(e) | e \in E(G_1) - E(G^{\triangleleft})\}$ is thin, since each time we chose some $dt(e)$ we specified a pi Π_0, such that no $\Pi' \preceq \Pi_0$ could meet $dt(e)$ (see Observations 2 to 5 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, $|G^{\triangleleft}| \approxeq |G_1| \approxeq |G^{\not\between}|$ (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 be $dt(e)$).

7.5 Cleaning up the Articulation Points

Keeping to our plan, we now rid the articulation points of unwanted edges. For all $i, j \in \mathbb{N}$, perform as many shortcuts at y^i_j (in $G^{\triangleleft}$) as possible, without shortcutting edges belonging to R^i or A^i; call the new edges *shortcutting edges* (recall that we have already defined another kind of shortcutting edges in Section 7.4). We are left with a muldigraph G^y, where each y^i_j is incident with at most one edge not in R^i (nothing needs to be done at articulation points of arcs, 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 $d(e)$ for each edge e in $E(G_2) - E(G^{\triangleleft})$ the two edges of $G^{\triangleleft}$ shortcutted to give e, we prove that $|G_2| \approxeq |G^{\triangleleft}|$.

For each edge $e = uv \in E(G_2) - E(G^y)$, either e is a bond, or it represents a bond of Π, where Π is either a Π^l_0 for some l, or the C_1 of some arc analysis. Let y^i_j be the articulation point where e was shortcutted, and suppose that $u = y^i_j$.

In the case that $\Pi = \Pi^l_0$ for some l, Π sends its anchor to RL^i_- (by the choice of the y^i_j). By the construction of $G^{\triangleleft}$, there is an $a(\Pi^l_0)$-u–path Q containing only vertices of $\Pi^i_{j-1}\Pi^i_j, \Pi^i_{j+1}$ and A^i. By Condition 3, there is in $G^{\triangleleft}$ a v-$a(\Pi^l_0)$–path P containing only vertices of RL^l_- and A^l. Since $a(\Pi^l_0)$ is not an articulation point, no edge in P or Q could have been shortcutted, thus we may choose $dt(e) := P \cup Q$ as a detour for e.

In the case that Π is an arc, 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 $dt(e)$.

Clearly, the set $\{d(e) | e \in E(G_2) - E(G^y)\}$ is thin, and by Lemma 8, $|G^y| \approxeq |G_2| \approxeq |G^{\triangleleft}|$.

We 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 $|G^y| \approxeq |G^{\triangleleft}| \approxeq |G^{\not\between}| \approxeq |G^{\between}|$, we have $|G^y| \approxeq$

$|G_3|$. We know that $G^{\between}$ is Eulerian, thus, by the definition of 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 $E(G^y) = T \triangle E(G^{\between})$, 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, proving that G^y is Eulerian.

7.6 The Hamiltonisation

By Theorem 4 we obtain an end-faithful Euler tour σ of G^y. Replace every shortcutting edge in σ by the two edges it shortcuts (this is done by modifying σ on the interval of S^1 mapped on the shortcutting edge, so that this interval is mapped continuously and bijectively on the two shortcutted edges), and 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 muldigraph resulting from $G^{\between}$ after doubling all single edges; σ' is clearly end-faithful. We will use Condition 2 to transform σ' to a Hamilton circle of G^2.

A *pass* of σ' through some vertex x, is a trail $uexe'v$ traversed by σ'. For every $x \in V - \{y^*\}$, let i be the index of x in $P(W(x))$, and shortcut the two edges of all passes of σ' through x that do not contain $e_i^-(W(x))$. Moreover, shortcut in σ' the two edges of all passes of σ' through y^* that do not contain e^*. We claim that we did not create any edge between two vertices that are not neighbours in G^2.

It suffices to show that no edge has been shortcutted at both its endvertices. Suppose that $e = xv$ has been shortcutted at x. If $x = y^*$, then easily $e = e_1(W)$, where $W = W(v)$, thus e has not been shortcutted at $v = x_1(W)$. If $x \neq y^*$, then let $W = W(x)$ and suppose that $x = x_i(W)$. Again we will show that e has not been shortcutted at v.

If e lies in $P(W)$, then $e \neq e_i^-$ because e has been shortcutted at x. 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 $e_{i+1}^{\cap}$ belongs to the same pass of σ' through x_i that contains e_i^-. If $e = e_i^{\cap}$, then by the same argument, it belongs to the same pass of σ' through x_{i-1} that contains e_{i-1}^-, and has thus not been shortcutted at x_{i-1}. If $e = e_{i+1}^-$, again e cannot have been shortcutted at x_{i+1}, unless x_{i+1} is the pseudo-mouth of W; but if $x_{i+1} = v$ is the pseude-mouth of W, then e, e_i^- where represented in $G^{\vartriangleleft}$, so they belong to the same pass of σ' through x, which implies that e cannot have been shortcutted at x.

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

Thus no edge has been shortcutted at both its endvertices. Now for every couple of edges e, f that were shortcutted to give an edge h, change σ' on the interval I of S^1 previously mapped on $e \cup f$, so that it continuously maps I to h. This transforms σ' to a mapping $\tau : S^1 \to |G^2|$. As no edge $e = y^*v$ could have been shortcutted at v (see the beginning of the proof that no edge was shortcutted at both endvertices), $\tau(S^1)$ contains e^*, and the other edge in $\tau(S^1)$ incident with y^* is also in $E(G)$. By construction, τ traverses each vertex in V exactly once. By Lemma 8 we easily have $|G^2| \approxeq |G|$, and since $|G| \approxeq G^\sharp \approxeq |G^{\between\between}|$, it follows that τ is continuous and end-faithful. Thus τ is a Hamilton circle of G^2. This completes the proof of Theorem 8, which implies Theorem 3.

The fact that the square of a 2-connected finite graph G is Hamiltonian connected ([13]), also generalises to locally finite graphs:

Corollary 1. *The square of a 2-connected locally finite graph G is Hamiltonian connected, that is, for each pair of vertices x, y of G, there is a homeomorphic image in $|G^2|$ of the unit interval with endpoints x, y.*

Proof. Add a new vertex y^* to G, join it to x, y with edges and apply Theorem 8. □

8 Final Remarks

In this paper we generalised Fleischner's Theorem to locally finite graphs. What about generalising other sufficient conditions for the existence of a Hamilton cycle? In general, as in our case, it is a hard task, and it is not clear why it should be possible. See for example [2], where Tutte's Theorem [16], that a finite 4-connected planar graph has a Hamilton cycle, is partially 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 some suitable Euler tour in some auxiliary graph, and then try to generalise the proof to the infinite case using Theorem 4.

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

Corollary 2. *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 end-faithful Euler tour of G, that 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 injectively onto 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 end 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. $\square$

Zhan [17] proved that every finite 7-connected line graph is hamiltonian. In view of Corollary 2, a generalisation to locally finite graphs looks plausible:

Conjecture 1. *Every locally finite* 7*-connected line graph has a Hamilton circle.*

References

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

[2] H. Bruhn and X. Yu. Hamilton cycles in infinite planar graphs: see Chapter 6 of H. Bruhn, Infinite circuits in locally finite graphs, PhD thesis, Universität Hamburg, 2005.

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

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

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

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

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

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

[9] A. Georgakopoulos. Connected but not path-connected subspaces of infinite graphs. To appear in Combinatorica.

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

[11] R. Halin. A note on Menger's theorem for infinite locally finite graphs. *Abh. Math. Sem. Univ. Hamburg*, 40:111–114, 1974.

[12] P. Heinrich. Abzählbar-unendliche graphen mit hamiltonscher dritter potenz. *Math. Nachr.*, 82:231–239, 1978.

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

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

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

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

[17] S. Zhan. On hamiltonian line graphs and connectivity. *Discrete Mathematics*, 89:89–95, 1991.