Connected but not path-connected subspaces of infinite graphs

Angelos Georgakopoulos

Abstract

Solving a problem of Diestel [9] relevant to the theory of cycle spaces of infinite graphs, we show that the Freudenthal compactification of a locally finite graph can have connected subsets that are not path-connected. However we prove that connectedness and path-connectedness do coincide for all but a few sets, which have a complicated structure.

1 Introduction

The Freudenthal compactification $|G|$ of a locally finite graph G is a well-studied space with several applications. For example, Cartwright, Soardi and Woess [8] study it in the context of random walks on infinite graphs and show that it coincides with the Martin compactification whenever G is a tree (see also Woess [21]). Polat [18] investigates its subspace topology on the set of vertices and ends, and relates the existence of certain spanning trees to the metrizability of this space. Bruhn, Diestel, Kühn and Stein [2-6,9,11-14] use it in order to define topological notions of paths, cycles and spanning trees that permit the extension of classical theorems about the cycle space of finite graphs to infinite ones. Finally, Stein [19, 20] and Bruhn and Yu [7] have begun to tackle topological variants of extremal-type problems in $|G|$, such as hamiltonicity or forcing highly connected subgraphs, that are standard for finite graphs but would otherwise have no counterpart for infinite graphs.

However, the following fundamental problem has remained unsolved:

Problem 1 ([9]). *Is every connected subspace of $|G|$ path-connected?*

Apart from being interesting as a basic topological question in its own right, this problem is also important from the graph-theoretical point of view. Indeed, the basic concepts in the context of [2-6,9,11-14] are circles and topological spanning trees, generalising cycles and spanning trees respectively. In order to prove that a certain subspace of $|G|$ is a circle or a topological spanning tree, one has to show that it is path-connected, but it is often much easier to show that it is (topologically) connected. See for example Theorem 8.5.8 in [10], which summarizes the basic properties of the cycle space of a locally finite graph. The reduction of path-connectedness to connectedness simplifies its proof considerably in comparison to the original proof in [13, Theorem 5.2]. Lemma 8.5.13 in [10] is an example of how reducing path-connectedness to connectedness can facilitate proving the existence of a topological spanning tree, which can otherwise be a tedious task as witnessed by the proof of Theorem 5.2 in [14]. Further examples are given in Exercises 65 and 70 of [10], which describe some fundamental properties of circles and topological spanning trees: their proofs become easy when the path-connectedness required is replaced with connectedness, while without this tool they would be arduous and long.

Diestel and Kühn [14] have shown that every closed connected subspace of $|G|$ is path-connected, and expressed a belief that the answer to Problem 1 should be positive also in general. However, we shall construct a counterexample (Section 3):

Theorem 1. *There exists a locally finite graph G such that $|G|$ has a connected subspace which is not path-connected.*

While the construction for the proof of Theorem 1 is our main result, we also prove that 'most' connected subsets as above are indeed path-connected (Section 4):

Theorem 2. *Given any locally finite connected graph G, a connected subspace X of $|G|$ is path-connected unless it satisfies the following assertions:*

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

Since the existence of a connected but not path-connected subspace was rather unexpected, I consider it as the main result of this paper and present it first. The counterexample can well be read by itself, but it may look somewhat surprising. However, the proof of Theorem 2 will make it less surprising with hindsight: it will show why the example had to be the way it is.

2 Definitions

We are using the terminology of [10] for graph theoretical concepts and that of [1] for topological ones.

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

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

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

[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.

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

It can be proved (see [11]) that in fact $|G|$ is the Freudenthal compactification [15] of the 1-complex G.

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

For any vertex $u \in V$ let $N^i(u)$ denote the set of vertices of G whose distance from u is at most i (including u).

A continuous map from the real unit interval $[0,1]$ to a topological space X is a (topological) *path* in X. A homeomorphic image (in the subspace topology) of $[0,1]$ in a topological space X will be called an *arc* in X. We will need the following lemma from elementary topology [16, p. 208]:

Lemma 1. *The image of every path with distinct endpoints x,y in a Hausdorff space X contains an arc in X between x and y.*

3 The counterexample: proof of Theorem 1

In this section we prove Theorem 1. Let $G = (V,E)$ be a graph. A subgraph consisting of a path xyz of order 3 and three disjoint rays starting at x,y,z respectively will be called a *trident*. The path xyz is the *spine* of the trident, and the rays are its *spikes*. The ends of G that contain the rays of the trident will be called, with slight abuse of terminology, the *ends of* the trident.

We will now recursively construct an infinite locally finite graph G and a subgraph X^*, which will be a collection of disjoint double rays of G. At the same time we will define a sequence of trees $\{T_i\}_{i<\omega}$ of auxiliary use. All vertices of any T_i, apart from their common root r, will be tridents in G.

Start with two tridents t_0, t_1 with a common spine, but otherwise disjoint (Figure 2). Put the three disjoint double rays formed by their spikes in X^*. Let T_0 consist of its root r and t_0, t_1 each joined to r.

Now perform ω steps of the following type. At step i, consider separately every trident v in G that is a leaf of T_i. Denote the spikes of v as α, β, γ and add to G three disjoint double rays and 6 further edges as in Figure 1 (these 6 edges are shown in thin continuous lines) to obtain the three new tridents with spikes μ, α, ν, and κ, β, λ, and o, γ, ξ. Add these tridents to T_i as neighbours of v. Let T_{i+1} be the tree resulting from such addition of three new tridents at every leaf of T_i; then T_{i+1} has no leaves in common with T_i. For every leaf of T_i, add to X^* the three double rays $\mu e_{\mu\xi} \xi$, $\nu e_{\nu\kappa} \kappa$ and $\lambda e_{\lambda o} o$ shown in dashed lines in Figure 1. (Note that the spikes α, β, γ of the old trident each contain a spike of one of the new tridents. Thus each ray will eventually participate in an infinite number of tridents.) Figure 2 shows the graph after the first and part of the second step.

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

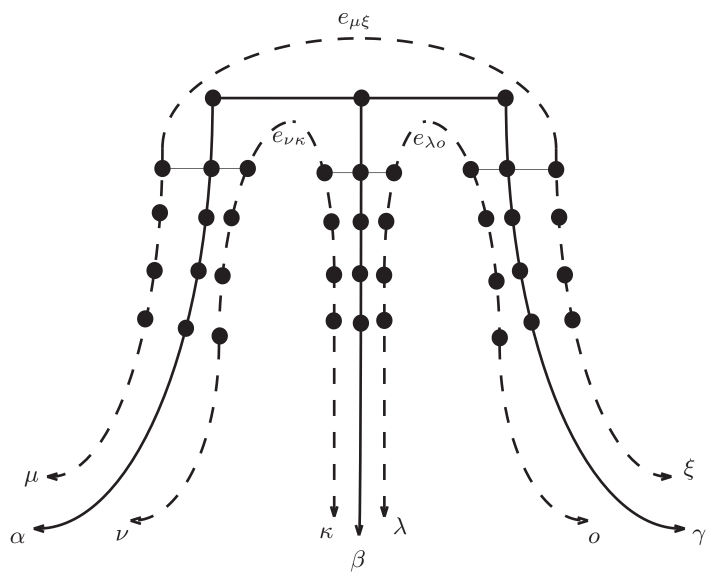

Figure 1: Three new tridents, with spikes μ, α, ν, and κ, β, λ, and o, γ, ξ

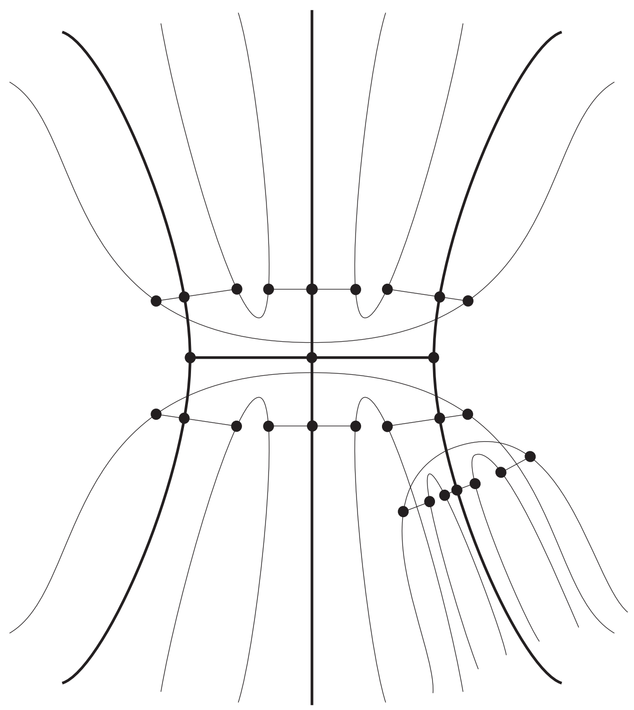

Figure 2: The first steps of the construction of G. The thick lines depict t_0 and t_1.

We will construct a connected set $X \subset |G|$ that is not path-connected. The path-components of X containing vertices, demanded by Theorem 2, will be the closures of the double rays of X^*. In order to supply the singleton ends, we will now divide the implicit ends between X and its complement in $|G|$, in such a way that

$$\text{neither } X \text{ nor } \Omega \backslash X \text{ contains a closed set of continuum many ends.} \qquad (*)$$

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

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

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

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

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

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

If t is any T-trident with spine xyz and α one of its ends, then $U(t) := O(\{x, y, z\}, \alpha)$ is a basic open set that does not depend on the choice of α; note that, by virtue of the '6 additional edges' of Figure 1, all three spikes of t have a subray in the same component of $G - \{x, y, z\}$. Then $\mathcal{U} := \{U(t) | t \text{ is a } T\text{-trident}\}$ is a basis of the open neighbourhoods of the ends of G, because for every end and every finite $S \subset V(G)$ there is a $U(t)$ that contains the end and misses S.

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

Next, we show that if a T-trident t is two-coloured, then there are two-coloured T-tridents r, s such that $U(r), U(s)$ are disjoint proper subsets of $U(t)$ (In other words r and s are both descendants of t, but not of each other). Let the tridents x, y, z be the children of t in T. We may assume that the spike of t that meets y is green, while its other two spikes are red (Figure 3).

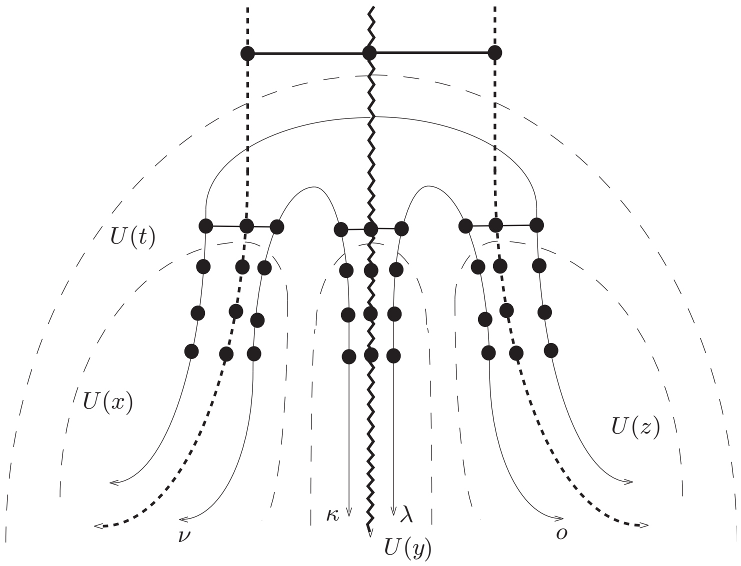

Figure 3: $U(t)$ and its subneighbourhoods.

Now consider the three thin double rays in Figure 3. If any of these is green, then at least two of the tridents x, y, z will be two-coloured. So let us assume that all those three double rays are red. But now y is coloured like t (one spike green, the other two red), and we may repeat the argument with y in the place of t. We continue recursively to find a descending ray $y_0y_1y_2\dots$ in T (with $y_0 = t$ and $y_1 = y$) of two-coloured tridents. But the sets $U(y_i)$ form a neighbourhood basis of the end ω of the green spike of t. This contradicts the fact that $\omega \in O_g$ and O_g is open.

We have thus shown that T contains a subdivision B of the infinite binary tree all whose branch vertices are 2-coloured tridents. Let $\sigma = x_1x_2\dots$ be any descending sequence of branch vertices of B. Then $\bigcap_{i\in\mathbb{N}} U(x_i)$ contains a unique end, $\omega(\sigma)$. As the $U(x_i)$ form a neighbourhood basis of $\omega(\sigma)$ and are all 2-coloured, so $\sigma \in \Omega\backslash(O_r \cup O_g)$. Since B contains continuum many such sequences σ, and their corresponding ends $\omega(\sigma)$ are clearly distinct, the set Ω' of all these ends $\omega(\sigma)$ is a subset of $\Omega\backslash(O_r \cup O_g)$ containing continuum many ends. As $O_r \cup O_g$ is open, the closure of Ω' still lies in $\Omega\backslash(O_r \cup O_g) \subseteq \Omega\backslash X$. This contradicts $(*)$, and completes our proof that X is connected.

It remains to prove that X is not path-connected. Suppose it is. Then any two distinct implicit ends $x, y \in X$ are connected by a path in X, and by Lemma 1 there is also an x–y arc A in X. We show that A contains continuum many ends, which will contradict $(*)$.

It is easy to confirm (see Lemma 4 below) that A must contain a vertex of X^*. Clearly, the double ray $R \in X^*$ containing this vertex is a subarc of A. Let A' and A'' denote the path-components of $A\backslash R$, which are subarcs of A preceding and following R. As before, A' and A'' each contain a double ray from

X^*, R' and R'' say. These double rays cannot share an end with R, because by construction no end contains more than one ray of X^*, hence R' and R'' split A' and A'' in two smaller subarcs.

Repeating recursively on each subarc of the previous step, we see that A contains a set $\mathcal{R}$ of infinitely many double rays, arranged like the segments of the unit interval removed to form the Cantor set. Imitating the corresponding proof, we see that A contains a set C of continuum many points that are limits of the ends of the double rays in $\mathcal{R}$. But only ends can be limits of ends, so C is a set of ends of X.

The arc A is closed because it is compact (as image of the compact space $[0, 1]$) and $|G|$ is a Hausdorf space (see [12] for a proof of this fact). The set of ends that lie on A is also closed, because its complement in $|G|$ consists of the complement of A in $|G|$ plus a set of vertex and edge points, and each of the later has an open neighbourhood that avoids all ends. Since this set contains C, it follows that A contains a closed subset of $\Omega \cap X$ with continuum many elements contradicting $(*)$.

This completes the proof that X is not path-connected and hence the proof of Theorem 1.

4 Positive results

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

The aim of this section is to prove Theorem 2. To this end we will first have to develop some intermediate results.

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

Lemma 2. *For every point $x \in X \backslash \Omega$ there is an open neighbourhood $U = U(x)$ of x such that $U \cap X \subseteq c(x)$.*

Proof.

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

$$\textit{one of the closed intervals } [u, x],\ [x, v] \textit{ lies in } X \textit{ as well.} \tag{2}$$

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

Thus (2) holds and we may assume without loss of generality that $[u, x] \subset X$. Now if X contains an interval $(x, w) \subset [x, v]$ we can set $U(x) = (u, w)$. Otherwise there is a point $v' \in (x, v)$ such that $(x, v') \cap X = \emptyset$, and we can set $U(x) = (u, v')$. For if no such v' exists, then there are points of X on (x, v) arbitrarily close to x. But for every such point y we can prove that $[y, v] \subset X$ the same way we proved (2) ($[u, y] \not\subset X$ because then we would have the previous case) and thus $[x, v] \subset X$ contradicting the assumption that X contains no interval $(x, w) \subset [x, v]$.

Now assume x is a vertex. By a similar argument as above we can prove that for every edge $xv \in E(G)$ there is a point $v' \in (x, v)$ such that either $(x, v') \subset X$

or $(x, v') \cap X = \emptyset$. Let S be a set that contains one such point for each edge incident with x. Then we can set $U(x) = \bigcup_{v' \in S} [x, v')$. □

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

Lemma 3. *For every end $\omega \in X$ and every path-component $c' \neq c(\omega)$ of X there is an open neighbourhood $U = U(c', \omega)$ of ω such that $U \cap c' = \emptyset$.*

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

The following will be needed for the proof of Lemma 3. We will say that a topological path *traverses* an edge xy if it contains $[x, y]$ as a subpath.

Lemma 4. *Any topological path that connects some point of a basic open neighbourhood U of an end to a point outside U must traverse some edge xy with $x \in U, y \notin U$.*

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

Proof of Lemma 3. Suppose there is a path-component c' of X and an end $\omega \in X \backslash c'$ every open neighbourhood of which meets c'. By Lemma 4, c' must contain a vertex u.

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

Define $M_i = (S_{i+1} \backslash S_i) \cap c' \cap U_i$, for all $i \geq 0$ (Figure 4). Each M_i is a set of candidates for the vertex a_i mentioned above. Instead of choosing them arbitrarily, we will make use of Konig's infinity lemma [17] to find a sequence of appropriate a_i.

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

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

For $i > 0$ pick any point of $U_i \cap c'$ and any topological path from that point to u. By Lemma 4, and since $u \notin U_i$, this path traverses one of the edges between a vertex w in U_i and a vertex outside it. By definition, M_i contains this vertex w.

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

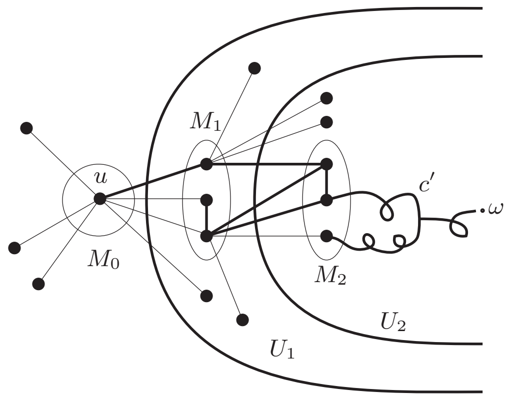

Figure 4: U_i and M_i.

a first point y in M_{i-1}. By Lemma 4, the subpath from x to y lies in U_{i-1}, so xy is an edge of $\mathcal{G}$.

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

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

□

As a consequence of Lemmas 2 and 3 we have the following:

Corollary 5. *The path-components of X are closed in its subspace topology.*

This implies that any counterexample to Problem 1 must contain infinitely many path-components. In fact we can prove something stronger:

Lemma 6. *Every connected but not path-connected $X \subseteq |G|$ contains uncountably many path-components.*

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

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

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

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

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

Then for every $i < \omega$ perform the following actions (see Figures 5 and 6):

1. Declare live all the basic open sets of the form $O(N^i(u), \omega)$, with $\omega \in \Omega \cap X$ that lie in live open sets of the previous step (note that $O(N^i(u), \omega) \subsetneq O(N^{i-1}(u), \omega)$).

2. Colour any still uncoloured path-component c that meets more than one live open set with the colour of the boss of the parent open set, i.e. the live open set of the previous step in which c lies (it must lie in one, because if it met more than one of them it would have been coloured in a previous step). Note that there are only finitely many such path-components in any step, because by Lemma 4 each of them must contain an edge that crosses some basic open set and there are only finitely many such edges. Finally give the newly coloured path-components the next free handicaps, one to each.

3. If a live basic open set does not meet any green path-components, then colour all path-components that lie in it red, put it in O_r and declare it not live. Proceed similarly with colours switched and O_g instead of O_r.

4. For every live basic open set, let the path-component of smallest handicap that meets it be its boss.

5. If c_i is still uncoloured, give it the colour of the boss of the live set in which it lies (it lies in one since it is still uncoloured) and the next free handicap.

We claim that after this process every end of X is put in either O_r or O_g. Indeed, because of action 5, for every end e of X, $c(e)$ gets a colour and a handicap sometime. By Lemma 3 and the fact that there are only finitely many path-components of smaller handicap, at some step j, e will lie in a live basic open set U that avoids all path-components of smaller handicap ($N^i(u)$ contains any finite vertex set for i large enough, if we assume, without loss of generality, that G is connected).

At this point, U only meets finitely many coloured path-components (see comment in action 2). In the steps following step j, e's path-component will

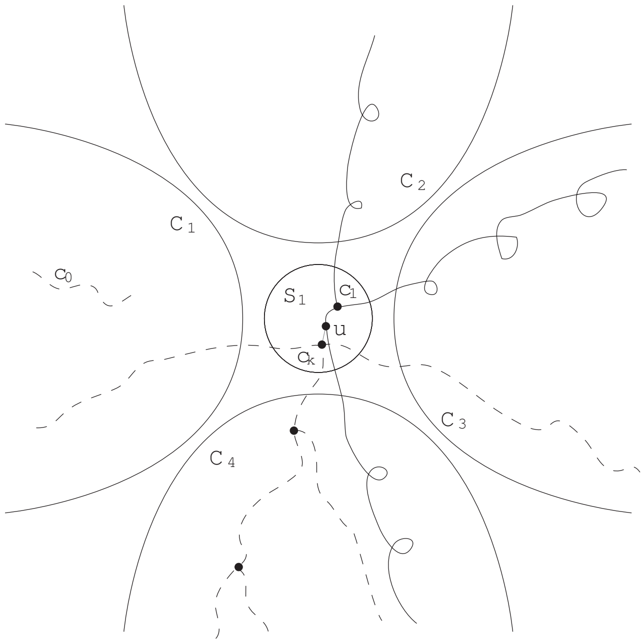

Figure 5: Possible colourings after step 1. Dashed lines depict red path-components and continous lines green ones. The path-component c_k meets several live sets, so it took the colour of c_0, the boss of $|G|$. The basic open set C_1 will be put in O_r and C_2 will be put in O_g; then they will be declared not live. The boss of both C_3 and C_4 is c_1.

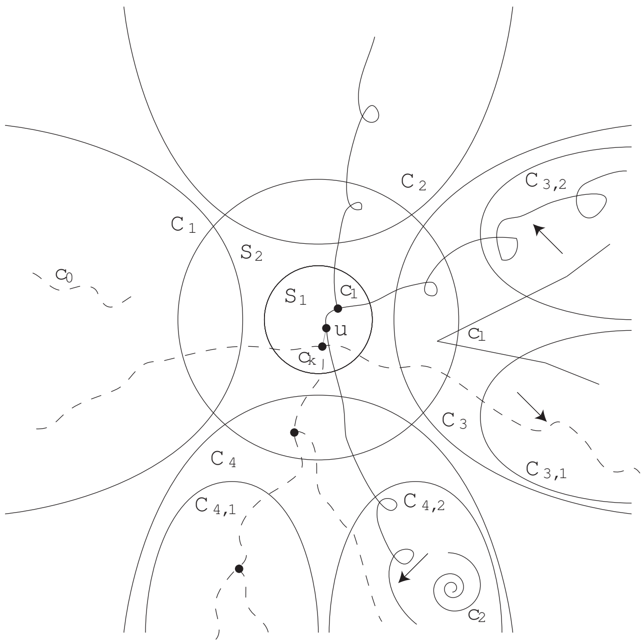

Figure 6: Possible colourings after step 2. $C_{4,1}$ will be put in O_r and $C_{3,2}$ in O_g. The path-component c_2 received the colour of c_1, the boss of $C_{4,2}$, and c_l received the colour of the boss of C_3, again c_1. The arrows show the bosses of the open sets that are live after the completion of this step.

always be the boss of the current live open set in which e lies (action 4) and thus no path-component that meets such a set will be coloured with the opposite colour (action 2).

Again by Lemma 3, e will at some later step lie in a basic open set U' that avoids all path-components of the opposite colour that met U at step j. This U' thus meets only the colour of e, so it will be classified in one of O_r, O_g.

Thus our claim is true and we have divided $X \cap \Omega$ into two open sets whose intersection does not meet X. Now for each vertex or edge point of X find a basic open set, supplied by Lemma 2, that avoids all other path-components, and put it in O_r if the point belongs to a red path-component, or in O_g if it belongs to a green one. Since $X \subseteq O_r \cup O_g$, and $O_r \cap O_g$ does indeed not meet X, the connectedness of X is contradicted.

□

Using the above Lemmas we can now prove Theorem 2:

Proof of Theorem 2. Suppose X is not path-connected. Since G is locally finite and connected, X has only countably many path-components containing vertices, so by Lemmas 6 and 4 there must be uncountably many path-components that are singleton ends.

If $c_1, c_2, \ldots, c_n$ are the only path-components of X that contain vertices, then pick a singleton end $\omega \in X$, and for each i an open neighbourhood O_i of ω that avoids c_i, supplied by Lemma 3. Let $O = O(S, \omega)$ be a basic open neighbourhood of ω contained in $\bigcap_{i \leq n} O_i$ (if $n = 0$ then let O be any basic open neighbourhood of ω that avoids at least one end of X). Every point of $X \backslash O$ has an open neighbourhood that does not meet O: for vertices and edge points this open neighbourhood is supplied by Lemma 2 and for an end $e \notin O$ the neighbourhood $O = O(S, e)$ does indeed not meet O. Thus $X \cap O$ is open and closed in the subspace topology of X, a contradiction since X is connected. This proves that X must have infinitely many path-components that contain vertices.

Finally, let us show that every path-component of X contains ends. By Lemma 2, a path-component containing no end is open. Since it is also closed (Corollary 5), the connectedness of X is contradicted if such a path-component exists.

□

References

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

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

[3] H. Bruhn and R. Diestel. Duality in infinite graphs. To appear in *Comb., Probab. Comput.*

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

[5] H. Bruhn and M. Stein. MacLane's planarity criterion for locally finite graphs. To appear in *J. Combin. Theory (Series B)*.

[6] H. Bruhn and M. Stein. On end degrees and infinite circuits in locally finite graphs. Preprint 2004.

[7] 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.

[8] D.I. Cartwright, P.M. Soardi, and W. Woess. Martin and end compactifications for nonlocally finite graphs. *Trans. Am. Math. Soc.*, 338:679–693, 1993.

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

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

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

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

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

[14] R. Diestel and D. Kühn. Topological paths, cycles and spanning trees in infinite graphs. 2004.

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

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

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

[18] N. Polat. Ends and multi-endings I. *J. Combin. Theory (Series B)*, 67:86–110, 1996.

[19] M. Stein. Arboriticity and tree-packing in locally finite graphs. To appear in *J. Combin. Theory (Series B)*.

[20] M. Stein. Forcing highly connected subgraphs in locally finite graphs. Preprint 2005.

[21] W. Woess. *Random walks on infinite graphs and groups.* Cambridge University Press, 2002.