

Dual trees must share their ends

Reinhard Diestel

Julian Pott

Abstract

We extend to infinite graphs the matroidal characterization of finite graph duality, that two graphs are dual iff they have complementary spanning trees in some common edge set. The naive infinite analogue of this fails.

The key in an infinite setting is that dual trees must share between them not only the edges of their host graphs but also their ends: the statement that a set of edges is acyclic and connects all the vertices in one of the graphs iff the remaining edges do the same in its dual will hold only once each of the two graphs' common ends has been assigned to one graph but not the other, and 'cycle' and 'connected' are interpreted topologically in the space containing the respective edges and precisely the ends thus assigned.

This property characterizes graph duality: if, conversely, the spanning trees of two infinite graphs are complementary in this end-sharing way, the graphs form a dual pair.

1 Introduction

It is well known (and easy to show) that two finite graphs are dual, in the usual sense that the circuits of one are the bonds of the other [8], if and only if they can be drawn with a common abstract set of edges so that the edge sets of the spanning trees of one are the complements of the edge sets of the spanning trees of the other:

Theorem 1. *Let $G = (V, E)$ and $G^* = (V^*, E)$ be connected finite graphs with the same abstract edge set. Then the following statements are equivalent:*

- (i) G and G^* are duals of each other.
- (ii) *Given any set $F \subseteq E$, the graph (V, F) is a tree if and only if (V^*, F^c) is a tree.*

For infinite dual graphs G and G^* (see [4]), Theorem 1 (ii) will usually fail: when (V, F) is a spanning tree of G , the subgraph (V^*, F^c) of G^* will be acyclic but may be disconnected. For example, consider as G the infinite $\mathbb{Z} \times \mathbb{Z}$ grid, and let F be the edge set of any spanning tree containing a two-way infinite path, a *double ray* R . Then the edges of R will form a cut in G^* , so (V^*, F^c) will be disconnected.

Although the graphs (V^*, F^{\complement}) in this example will always be disconnected, they become arc-connected (but remain *acirclic*) when we consider them as closed subspaces of the topological space obtained from G^* by adding its end. Such subspaces are called *topological spanning trees*; they provide the ‘correct’ analogues in infinite graphs of spanning trees in finite graphs for numerous problems, and have been studied extensively [9, 10]. For $G = \mathbb{Z} \times \mathbb{Z}$, then, the complements of the edge sets of ordinary spanning trees of G form topological spanning trees in G^* , and vice versa (as $\mathbb{Z} \times \mathbb{Z}$ is self-dual).

It was shown recently in the context of infinite matroids [5] that this curious phenomenon is not specific to this example but occurs for all dual pairs of graphs: neither ordinary nor topological spanning trees permit, by themselves, an extension of Theorem 1 to infinite graphs, but as soon as one notion is used for G and the other for G^* , the theorem does extend. The purpose of this paper is to explain this seemingly odd phenomenon by a more general duality for graphs with ends, in which it appears as merely a pair of extreme cases.

It was shown in [6] that 2-connected dual graphs do not only have the ‘same’ edges but also the ‘same’ ends: there is a bijection between their ends that commutes with the bijection between their edges so as to preserve convergence of edges to ends.¹ Now if G and G^* are dual 2-connected graphs with edge sets E and end sets Ω , our result is that if we specify *any* subset Ψ of Ω and consider topological spanning trees of G in the space obtained from G by adding only the ends in Ψ , then Theorem 1 (ii) will hold if the subgraphs (V^*, F^{\complement}) of G^* are furnished with precisely the ends in $\Omega \setminus \Psi$. (Our earlier example is the special case of this result with either $\Psi = \emptyset$ or $\Psi = \Omega$.) And conversely, if the spanning trees of two graphs G and G^* with common edge and end sets complement each other in this way for some—equivalently, for every—subset Ψ of their ends then G and G^* form a dual pair.

Here, then, is the formal statement of our theorem. A graph G is *finitely separable* if any two vertices can be separated by finitely many edges; as noted by Thomassen [13, 14], this slight weakening of local finiteness is necessary for any kind of graph duality to be possible. The Ψ -trees in G , for subsets Ψ of its ends, will be defined in Section 2. Informally, they are the subgraphs that induce no cycle or topological circle in the space which G forms with the ends in Ψ (but no other ends) and connect any two vertices by an arc in this space.

Theorem 2. *Let $G = (V, E, \Omega)$ and $G^* = (V^*, E, \Omega)$ be finitely separable 2-connected graphs with the same edge set E and the same end set Ω . Then the following assertions are equivalent:*

- (i) G and G^* are duals of each other.
- (ii) For all $\Psi \subseteq \Omega$ and $F \subseteq E$ the following holds: F is the edge set of a Ψ -tree in G if and only if F^{\complement} is the edge set of a Ψ^{\complement} -tree in G^* .
- (iii) There exists a set $\Psi \subseteq \Omega$ such that for every $F \subseteq E$ the following holds:

¹See the end of Section 2 for a more formal definition.

F is the edge set of a Ψ -tree in G if and only if $F^{\mathbb{G}}$ is the edge set of a $\Psi^{\mathbb{G}}$ -tree in G^* .

Setting $\Psi = \emptyset$ in (ii) and (iii) as needed, we reobtain the following result from [5]:

Corollary 3. *Two 2-connected and finitely separable graphs $G = (V, E, \Omega)$ and $G^* = (V^*, E, \Omega)$ are dual if and only if the following assertions are equivalent for every $F \subseteq E$:*

- (i) F is the edge set of a spanning tree of G ;
- (ii) $F^{\mathbb{G}}$ is the edge set of a topological spanning tree of G^* . □

We shall prove Theorem 2, extended by another pair of equivalent conditions in terms of circuits and bonds, in Sections 3–4.

2 Definitions and basic facts

All the graphs we consider in this paper will be *finitely separable*, that is, any two vertices can be separated by finitely many edges.

We think of a *graph* as a triple (V, E, Ω) of disjoint sets, of *vertices*, *edges*, and *ends*, together with a map $E \rightarrow V \cup [V]^2$ assigning to every edge either one or two vertices, its *endvertices*, and another map mapping the ends bijectively to the equivalence classes of *rays* in the graph, its 1-way infinite paths, where two rays are *equivalent* if they cannot be separated by finitely many vertices. In particular, our ‘graphs’ may have multiple edges and loops. For the complement of F in E , and of Ψ in Ω , we write $F^{\mathbb{G}}$ and $\Psi^{\mathbb{G}}$, respectively.

Let $G = (V, E, \Omega)$ be a graph, and let X be the topological 1-complex formed by its vertices and edges. In X , every edge is a topological copy of $[0, 1]$ inheriting also its metric. We denote the topological interior of an edge e by \mathring{e} , and for a set $F \subseteq E$ of edges we write $\mathring{F} := \bigcup_{e \in F} \mathring{e}$.

Let us define a new topology on $X \cup \Omega$, to be called **VTOP**. We do this by specifying a neighbourhood basis for every point. For points $x \in X$ we declare as open the open ϵ -balls around x in X with $0 < \epsilon < \delta$, where δ is the distance from x to a closest vertex $v \neq x$. For points $\omega \in \Omega$, note that for every finite set $S \subseteq V$ there is a unique component $C = C(S, \omega)$ of $G - S$ that contains a ray from ω . Let $\hat{C} = \hat{C}(S, \omega) \subseteq X \cup \Omega$ be the set of all the vertices in C , all the interiors of edges incident with a vertex in C , and all the ends represented by a ray in C . We declare all these sets \hat{C} as open, thus obtaining for ω the neighbourhood basis

$$\{\hat{C}(S, \omega) \subseteq X \cup \Omega : S \subseteq V, |S| < \infty\}.$$

We write $|G|$ for the topological space on $X \cup \Omega$ endowed with this topology.² In topological contexts we shall also write G for the subspace $|G| \setminus \Omega$. (This has the same points as X , but a different topology unless G is locally finite.)

²This differs a little from the definition of $|G|$ in [8] when G is not locally finite.

If ω and S are as above, we say that S *separates* ω in G from all the ends that have no ray in $C(S, \omega)$ and from all vertices in $G - C(S, \omega) - S$. Distinct ends are, by definition of ‘distinct’, separated by some finite set of vertices.

Let us show that $|G|$ is first-countable. We shall need this only when G is 2-connected, but it holds in general (for our finitely separable G):

Lemma 4. *Every point of $|G|$ has a countable neighbourhood basis.*

Proof. For $x \in G$ this is clear from the definition of VTop.

For $x \in \Omega$ it is clear if G is countable, since in that case there are only finitely many finite vertex sets S , and the sets $\hat{C}(S, x)$ form a neighbourhood basis for x . If G is 2-connected, it is indeed countable. Indeed, if not it must have a vertex v of infinite degree, found easily in any spanning tree. Then v has uncountably many neighbours in $G - v$, which is still connected. The union of all the paths between these neighbours in some fixed spanning tree of $G - v$ has another vertex of uncountable degree, w say, and so G contains uncountably many independent v - w paths. This contradicts our assumption that G is finitely separable.

Assume now that G is not 2-connected. Suppose first that x has a ray R in some block B of G . Then R lies in an end ω of B , which has a countable neighbourhood basis in $|B|$ consisting of sets $\hat{C}_B(S, \omega)$ with $S \subseteq V(B)$ finite and $C_B(S, \omega) \subseteq B$. Given any basic open neighbourhood $\hat{C}(S', x)$ of x in $|G|$, add to $S' \cap V(B)$ the cutvertices of G that separate a vertex of S' outside B from B , to form a finite set S of vertices in B . Then $\hat{C}(S, x) \subseteq \hat{C}(S', x)$ by the choice of B . These countably many $\hat{C}(S, x)$ thus form a neighbourhood basis for x in $|G|$.

Suppose finally that x has no ray in any block of G . Then each of its rays contains infinitely many cutvertices v of G that lie on some ray of the block-cutvertex tree of the component of G in which x has its rays. The sets $\hat{C}(\{v\}, x)$ for these countably many v then form a neighbourhood basis for x in $|G|$. \square

A vertex v *dominates* an end ω if G contains infinitely many paths from v to some ray in ω that pairwise meet only in v . When this is the case we call v and ω *equivalent*, and write \sim for the equivalence relation on $V \cup \Omega$ which this generates. Note that since G is finitely separable, no two vertices will be equivalent under \sim : every non-singleton equivalence class consists of one vertex and all the ends it dominates. A vertex and an end it dominates have no disjoint neighbourhoods in $|G|$. But two ends always have disjoint neighbourhoods, even if they are dominated by the same vertex.

We shall often work in the quotient space $\tilde{G} := |G|/\sim$, whose topology we denote as ITOP. Given a point $x \in |G|$, we write $[x]$ for its equivalence class under \sim , the point of \tilde{G} it represents. As different vertices are never equivalent, the vertices of G determine distinct \sim -classes; we call these the *vertices* of \tilde{G} . For any edge e of G we refer to the set $\{[x] : x \in e\}$ as an *edge* of \tilde{G} , which we continue to denote by e .

By definition of the quotient topology, a set $X \subseteq \tilde{G}$ is open in \tilde{G} if and only if $\bigcup X$, the union of all the subsets $[x] \in X$ of $|G|$, is open in $|G|$. Put another

way, the open sets in \tilde{G} are precisely the images under the projection $x \mapsto [x]$ of the subsets O of $|G|$ that are open in the topology of $|G|$ and closed³ under \sim . When we later need to construct open neighbourhoods of points p in \tilde{G} , this is how we shall find them: by constructing a set $O \subseteq |G|$ with these two properties which, moreover, contains p as a subset.

Example 1. Let G be the graph obtained from disjoint rays $R_n = v_n^1 v_n^2 \dots$, one for every $n \in \mathbb{N}$, by adding a new vertex v adjacent to all the other vertices. Write ω_n for the end of G containing R_n . Then v dominates all these ends, and $[v] = \{v, \omega_1, \omega_2, \dots\}$ is the unique point in \tilde{G} whose equivalence class in $|G|$ is not a singleton.

A basic open neighbourhood of $[v]$, then, is projection under $x \mapsto [x]$ of a set $O \subseteq |G|$ that is the union of a topological open star $B(v, \epsilon)$ around v (v itself plus an open segment from every edge at v , all of some common length ϵ) and sets of the form $\hat{C}(S_n, \omega_n)$, one for every n , where the S_n are finite sets of vertices that may differ for distinct n .

Thus, $[v]$ lies in the closure of the edges of any one of the rays R_n , viewed as edges of \tilde{G} , since every subset of $|G|$ of the form $\hat{C}(S_n, \omega_n)$ contains an edge of R_n as a subset (indeed, all but finitely many). But $[v]$ does not lie in the closure of the set $\{v_n^n \mid n \in \mathbb{N}\}$ of vertices of \tilde{G} , since we could form an open neighbourhood of $[v]$ by choosing, for each n , the set S_n so as to include v_n^n , in which case neither $B(v, \epsilon)$ nor any of the sets $\hat{C}(S_n, \omega_n)$ would contain any v_j^i .

More generally, for any set $\Psi \subseteq \Omega$ of ends we shall consider the subspace

$$|G|_\Psi := |G| \setminus \Psi^{\mathbb{G}}$$

and its quotient space

$$\tilde{G}_\Psi := |G|_\Psi / \sim,$$

whose topology we denote by Ψ -TOP. For $\Psi = \Omega$ we reobtain the space $\tilde{G}_\Omega = \tilde{G}$ defined above. We write $[x]_\Psi$ for the equivalence class of x in $|G|_\Psi$.

As before, the *vertices* of \tilde{G}_Ψ are those of its elements that are represented by vertices of G . All other points of \tilde{G}_Ψ are singleton classes $\{x\}$, with x either an inner point of an edge or an undominated end in Ψ . We will not always distinguish $\{x\}$ from x in these cases, i.e., call these $\{x\}$ also *inner point of edges* or *ends* of \tilde{G}_Ψ .

Informally, we may think of \tilde{G}_Ψ as formed from G in three steps:

- add the undominated ends from Ψ as new points, and make their rays converge to them;
- make the rays from any dominated end in Ψ converge to their unique dominating vertex;
- let the rays of ends in $\Psi^{\mathbb{G}}$ go to infinity without converging to any point.

³Note the different uses of the word ‘closed’ here: one topological, the other referring to closure under the relation \sim . Thus, for O as above we require that $x \sim y \in O$ implies $x \in O$.

Note that if Ψ contains a dominated end then $|G|_\Psi$ will fail to be Hausdorff, and if $\Psi^{\mathcal{G}} \neq \emptyset$ then \tilde{G}_Ψ will fail to be compact. But $\tilde{G} = \tilde{G}_\Omega$ is compact if G is 2-connected [7], and we shall see that \tilde{G}_Ψ is always Hausdorff (Corollary 8). Indeed, the choice of \sim is motivated by the fact that it is the finest equivalence relation on $|G|$ that makes \tilde{G}_Ψ Hausdorff for all choices of Ψ .

The diagram in Figure 1 shows the relationship between the spaces just defined. The subspace inclusion $\iota: |G|_\Psi \rightarrow |G|$ and the quotient projections $\pi: |G| \rightarrow \tilde{G}$ and $\pi_\Psi: |G|_\Psi \rightarrow \tilde{G}_\Psi$ are canonical, and $\sigma_\Psi: \tilde{G}_\Psi \rightarrow \tilde{G}$ is defined so as to make the diagram commute: it sends an equivalence class $[x]_\Psi \in \tilde{G}_\Psi$ to the class $[x] \in \tilde{G}$ containing it.

$$\begin{array}{ccc} |G|_\Psi & \xrightarrow{\iota} & |G| \\ \pi_\Psi \downarrow & & \downarrow \pi \\ \tilde{G}_\Psi & \xrightarrow{\sigma_\Psi} & \tilde{G} \end{array}$$

Figure 1: Spaces with ends, and their quotient spaces

Since G is finitely separable and hence no end is dominated by more than one vertex, σ_Ψ is injective: $\sigma_\Psi([x]_\Psi) = [x] \in \tilde{G}$ is obtained from $[x]_\Psi$ simply by adding those ends of $\Psi^{\mathcal{G}}$ that are dominated by a vertex in $[x]_\Psi$. As $|G|_\Psi$ carries the subspace topology induced from $|G|$, it is also easy to check that σ_Ψ is continuous. Its inverse σ_Ψ^{-1} can fail to be continuous; see Example 3 below.

The subtle differences between $|G|_\Psi$ and \tilde{G}_Ψ will often be crucial in this paper. But when they are not, we may suppress them for simplicity of notation. For example, given a subgraph H of G we shall speak of the *closure of H in \tilde{G}_Ψ* and mean the obvious thing: the closure in \tilde{G}_Ψ of its subspace $\pi_\Psi(H')$, where H' is H viewed as a subspace of $|G|_\Psi \subseteq |G|$.

By a *circle* in a topological space X we mean a topological embedding $S^1 \rightarrow X$, or its image. Since circles are compact and \tilde{G} is Hausdorff, σ_Ψ maps circles in \tilde{G}_Ψ to circles in \tilde{G} . Conversely, circles in \tilde{G} that ‘use’ only ends in Ψ define circles in \tilde{G}_Ψ ; this will be made precise in Lemma 12. The set of all the edges contained in a given circle in \tilde{G}_Ψ will be called a Ψ -*circuit* of G ; for $\Psi = \Omega$ we just speak of *circuits* of G . We shall not consider ‘circuits’ of circles in $|G|$ or $|G|_\Psi$.

As with circles, we use the term *path* in topological contexts both for continuous maps from $[0, 1]$, not necessarily injective, and for their images. For example, if A and B are the images of paths $\varphi, \varphi': [0, 1] \rightarrow \tilde{G}$ with endpoints $x = \varphi(0)$ and $y = \varphi(1) = \varphi'(0)$ and $z = \varphi'(1)$, we write $xAyBz$ for the ‘ x - y path’ in \tilde{G} that is the image of the concatenation of the paths φ and φ' . Note that, since \tilde{G}_Ψ is Hausdorff, every path in \tilde{G}_Ψ between two points x and y contains an x - y arc [12, p. 208].

A subspace of \tilde{G}_Ψ that is the closure in \tilde{G}_Ψ of the union of all the vertices and edges it contains is a *standard subspace of \tilde{G}_Ψ* . Circles in \tilde{G}_Ψ are examples of standard subspaces; this was shown in [10] for \tilde{G} , and follows for arbitrary Ψ

from Lemma 7 below. A standard subspace of \tilde{G}_Ψ that contains no circle is a Ψ -forest of G . A Ψ -forest is *spanning* if it contains all the vertices of \tilde{G}_Ψ . Note that, being closed, it then also contains all the ends of \tilde{G}_Ψ . A spanning arc-connected Ψ -forest of G is a Ψ -tree of G .

Thus, the \emptyset -trees of G are precisely its (ordinary) spanning trees, while its Ω -trees are its *topological spanning trees*, the arc-connected standard subspaces of \tilde{G} that contain all the vertices of \tilde{G} but no topological circle.

Example 2. Let G be obtained from a double ray D by adding a vertex v adjacent to all of D . This graph G has two ends, ω and ψ say, both dominated by v . The closure in \tilde{G} of the edges of D is a circle containing the ‘vertex’ $[v] = \{v, \omega, \psi\}$ of \tilde{G} , even though v does not lie on D . However for $\Psi = \{\psi\}$ the closure in \tilde{G}_Ψ of the same set of edges is not a circle but homeomorphic to a half-open interval. It thus is a Ψ -tree, since it contains the ‘vertex’ $\{v, \psi\}$ of \tilde{G}_Ψ as well as all its other vertices. The closure of the edges of D in \tilde{G}_\emptyset , on the other hand, is a connected \emptyset -forest but not a \emptyset -tree, since v lies in none of its points. Figure 2 shows a Ψ -tree for each choice of Ψ in this example.




Figure 2: Ψ -trees for $\Psi = \{\psi\}$, $\Psi = \emptyset$ and $\Psi = \{\omega, \psi\}$

If G and G^* are graphs with the same edge set, and such that the bonds of G^* (its minimal cuts) are precisely the circuits of G , then G^* is called a *dual* of G . If the *finite* bonds of G^* are precisely the finite circuits of G , then G^* is a *finitary dual* of G . Clearly, duals are always finitary duals. For finitely separable graphs, as considered here, the converse is also true [4, Lemmas 4.7–4.9].

If G^* is a dual of G , then G is a dual of G^* [4, Theorem 3.4], and G has a dual if and only if it is planar [4]. Moreover, there is a bijection $\omega \mapsto \omega^*$ between the ends ω of G and the ends ω^* of G^* such that a sequence e_1, e_2, \dots of edges (of both G and G^*) converges to end ω in $|G|$ if and only if it converges to ω^* in G^* , for all ends ω of G and ω^* of G^* [6].⁴ This bijection is clearly unique, and so there is no need to distinguish notationally between ω and ω^* . We express this informally by saying that G and G^* *have the same end sets* (as well as the same edge set).

⁴Convergence, here, is convergence of point sets: e_1, e_2, \dots *converges* to ω in $|G|$ iff every neighbourhood of ω in $|G|$ contains all but finitely many e_n as subsets, and likewise for ω^* in $|G^*|$.

3 Lemmas

Our main aim in this section is to prove some fundamental lemmas about the spaces $|G|$, $|G|_\Psi$, \tilde{G} and \tilde{G}_Ψ defined in Section 2: about their topological properties, and about their relationship to each other. Throughout the section, let $G = (V, E, \Omega)$ be a fixed finitely separable graph, and $\Psi \subseteq \Omega$ a fixed set of ends.

Before we get to these topological fundamentals, let us show that Ψ -trees always exist, and prove an easy lemma about how they relate to finite circuits and bonds. As to the existence of Ψ -trees, we can even show that there are always rather special ones: Ψ -trees that are connected not only topologically through their ends, but also as graphs:

Lemma 5. *If G is connected, it has a spanning tree T whose closure in \tilde{G}_Ψ is a Ψ -tree.*

Proof. It was shown in [4, Thm. 6.3] that G has a spanning tree T whose closure \bar{T} in \tilde{G} contains no circle. Let \bar{T}_Ψ denote the closure of T in \tilde{G}_Ψ . Then $\sigma_\Psi(\bar{T}_\Psi) \subseteq \bar{T}$. Since circles in \tilde{G}_Ψ define circles in \tilde{G} (by composition with σ_Ψ), \bar{T}_Ψ contains no circle either.

For a proof that \bar{T}_Ψ is arc-connected it suffices to show that every undominated end $\psi \in \Psi$ contains a ray $R \subseteq T$: then the arc $\pi_\Psi(R) \subseteq \bar{T}_\Psi$ connects the end $\{\psi\} \in \bar{T}_\Psi$ to a vertex, while all the vertices of \bar{T}_Ψ are connected by T . Pick a ray $R' \in \psi$ in G , say $R' = v_0 v_1 \dots$. By the star-comb lemma [8, Lemma 8.2.2], the connected graph $\bigcup_{n \in \mathbb{N}} v_n T v_{n+1}$ contains a subdivided infinite star with leaves in R' or an infinite comb with teeth in R' . As ψ is not dominated, we must have a comb. The back $R \subseteq T$ of this comb is a ray equivalent to R' that hence lies in ψ .

Being acirclic, arc-connected and spanning, \bar{T}_Ψ is a Ψ -tree. □

Lemma 6. *Assume that G is connected, and let $F \subseteq E$ be a finite⁵ set of edges.*

- (i) *F is a circuit if and only if it is not contained in the edge set of any Ψ -tree and is minimal with this property.*
- (ii) *F is a bond if and only if it meets the edge set of every Ψ -tree and is minimal with this property.*

Proof. (i) Assume first that F is a circuit. Then F is not contained in any Ψ -tree; let us show that every proper subset of F is. We do this by showing the following more general fact:

Every finite set F' of edges not containing a circuit extends to a spanning tree of G whose closure in \tilde{G}_Ψ is a Ψ -tree. (1)

To prove (1), consider a spanning tree T of G whose closure in \tilde{G}_Ψ is a Ψ -tree (Lemma 5). Choose it with as many edges in F' as possible. Suppose it fails to contain an edge $f \in F'$. Adding f to T creates a cycle C in $T + f$, which by

⁵See Section 5 for what can be said when F is infinite.

assumption also contains an edge $e \notin F'$. As C is finite, it is easy to check that $T + f - e$ is another spanning tree whose closure is a Ψ -tree. This contradicts our choice of T . This completes the proof of (1) and the proof of the forward implication in (i).

For the backward implication of (i), suppose that F is not contained in any Ψ -tree. Then by (1) it contains a circuit. If, in addition, it is minimal with the first property, it will in fact be that circuit, since we could delete any other edge without making it extendable to a Ψ -tree.

(ii) If F is a cut, $F = E(V_1, V_2)$ say, then the closures of $G[V_1]$ and $G[V_2]$ in \tilde{G}_Ψ are disjoint open subsets of $\tilde{G}_\Psi \setminus \overset{\circ}{F}$, so this subspace cannot contain a Ψ -tree. Thus, F meets the edge set of every Ψ -tree.

If F is even a bond, then both V_1 and V_2 induce connected subgraphs. By Lemma 5, these have spanning trees T_i ($i = 1, 2$) whose closures in \tilde{G}_Ψ are arc-connected and contain no circle.⁶ For every edge $f \in F$, the closure \overline{T}_Ψ of $T := (T_1 \cup T_2) + f$ in \tilde{G}_Ψ is a Ψ -tree of G : it still contains no circle, because no arc in $\overline{T}_\Psi \setminus \overset{\circ}{f}$ can cross the finite cut F from which it contains no edge (as above). So F is minimal with the property of meeting the edge set of every Ψ -tree.

Conversely, let us assume that F meets the edge set of every Ψ -tree, and show that F contains a bond. Let T be a spanning tree of G whose closure in \tilde{G}_Ψ is a Ψ -tree (Lemma 5), chosen with as few edges in F as possible. By assumption, T has an edge f in F . If the bond B of G between the two components of $T - f$ contains an edge $e \notin F$, then $T - f + e$ is another spanning tree whose closure is a Ψ -tree (as before) that contradicts our choice of T . So B contains no such edge e but is contained in F .

If F is minimal with the property of containing an edge from every Ψ -tree, it must be equal to the bond it contains. For by the forward implication of (ii) already proved, any other edge could be deleted from F without spoiling its property of meeting the edge set of every Ψ -tree. \square

We begin our study of the spaces introduced in Section 2 by showing that finite separability extends from G to \tilde{G}_Ψ :

Lemma 7. *For every two points $p, q \in \tilde{G}_\Psi$ that are not inner points of edges there exists a finite set F of edges and disjoint open subsets O_p, O_q of \tilde{G}_Ψ such that $p \in O_p$ and $q \in O_q$ and $\tilde{G}_\Psi \setminus \overset{\circ}{F} \subseteq O_p \cup O_q$.*

Proof. Let us write $p = [x]_\Psi$ and $q = [y]_\Psi$, where x and y are either vertices or undominated ends of G . We shall find a finite cut F of G , with bipartition (X, Y) of V say, such that $x \in \overline{X}$ and $y \in \overline{Y}$, where \overline{X} and \overline{Y} denote closures of X and Y in $|G|_\Psi$. Since F is finite, $\overline{G[X]}$ and $\overline{G[Y]}$ (closures in $|G|_\Psi$) then partition $|G|_\Psi \setminus \overset{\circ}{F}$ into disjoint open sets that are closed under equivalence, so their projections under π_Ψ partition $\tilde{G}_\Psi \setminus \overset{\circ}{F}$ into disjoint open sets containing p and q , respectively. We obtain O_p and O_q from these by picking an inner point

⁶We are applying Lemma 5 in the subgraphs $G[V_i]$. But since F is finite, the spaces $\widetilde{G[V_i]}_{\Psi_i}$ are canonically embedded in \tilde{G}_Ψ .

z_f from every $f \in F$ and adding each of the two segments of $f \setminus \{z_f\}$ to the partition side in which it has a vertex.

If x and y are vertices, then F exists by our assumption that G is finitely separable. Suppose now that y is an end. Let us find a finite set $S \not\ni x$ of vertices that separates x from y in G . If x is another end, then S exists since $x \neq y$. If x is a vertex, pick a ray $R \in y$. If there is no S as desired, we can inductively find infinitely many independent x - R paths in G , contradicting the fact that y is undominated.

Having found S , consider the component $C := C(S, y)$ of $G - S$. For each $s \in S$ we can find a finite set $S_s \subseteq C$ of vertices separating s from y in the subgraph of G spanned by C and s , since otherwise s would dominate y (as before). Let $S' := \bigcup_{s \in S} S_s$; this is a finite set of vertices in C that separates all the vertices of S from y in G . Since G is finitely separable, there is a finite set F of edges separating S from S' in G . Choose F minimal. Then, assuming without loss of generality that G is connected, every component of $G - F$ meets exactly one of the sets S and S' . Let X be the set of vertices in components meeting S , and let Y be the set of vertices in components meeting S' . Then (X, Y) is a partition of G crossed by exactly the edges in F , and it is easy to check that F has the desired properties. \square

It was proved in [10], under a weaker assumption than finite separability (just strong enough that \tilde{G} can be defined without identifying distinct vertices) that \tilde{G} is Hausdorff. For finitely separable graphs, as considered here, the proof is much simpler and extends readily to \tilde{G}_Ψ :

Corollary 8. \tilde{G}_Ψ is Hausdorff.

Proof. Finding disjoint open neighbourhoods for distinct points $p, q \in \tilde{G}_\Psi$ is easy if one of them is an inner point of an edge. If not, take sets O_p, O_q from Lemma 7. \square

Our next aim is to select from the basic open neighbourhoods $\hat{C}(S, \omega) \setminus \Psi^{\mathcal{G}}$ in $|G|_\Psi$ of ends $\omega \in \Psi$ some ‘standard’ neighbourhoods that behave well under the projection π_Ψ and still form neighbourhood bases of these points ω . Ideally, we would like to find for every end $\omega \in \Psi$ a basis of open neighbourhoods that are each closed under \sim in $|G|_\Psi$. That will not be possible, since ω can be separated topologically in $|G|_\Psi$ from every other end, including the ends ω' to which it is equivalent. But we shall be able to find, for every end $\omega \in \Psi$, bases of open neighbourhoods of all the ends $\omega' \sim \omega$ (including $\omega' = \omega$) that are each closed under \sim for all points other than ω' itself. Then the union of all these neighbourhoods, one for every end $\omega' \sim \omega$, plus an open star neighbourhood of their common dominating vertex, will be closed under \sim , and will thus be the pre-image in $|G|_\Psi$ of an open neighbourhood of the point $\pi_\Psi(\omega) = [\omega]_\Psi \in \tilde{G}_\Psi$. These neighbourhoods of $[\omega]_\Psi$ will form a basis of open neighbourhoods for $[\omega]_\Psi$ in \tilde{G}_Ψ .

Given a bond $F = E(V_1, V_2)$ of G and an end $\omega \in \Psi$ that lies in the $|G|$ -closure of V_1 but not of V_2 , let

$$\hat{C}_\Psi(F, \omega) \subseteq |G|_\Psi$$

denote the union of the $|G|_\Psi$ -closure of $G[V_1]$ with \hat{F} . For every vertex $v \in V_2$ we also call F a v - ω bond. Note that $\hat{C}_\Psi(F, \omega)$ depends only on F and ω : since F is a bond, $G - F$ has only two components, so V_1 and V_2 can be recovered from F and ω . Note also that every ray in ω has a tail in $\hat{C}_\Psi(F, \omega)$, so if it starts at v it must have an edge in F .

If $v \in V_2$ is an endvertex of all but finitely many of the edges in F , we say that F is v -cofinite. Then the set S of endvertices of F in V_2 is finite and separates ω from $V_2 \setminus S$.

Lemma 9. *Let $\omega \in \Psi$ be an end, and $v \in V$ a vertex.*

- (i) *If ω is undominated, then the sets $\{\hat{C}_\Psi(F, \omega) \mid F \text{ is a finite bond of } G\}$ form a basis of open neighbourhoods of ω in $|G|_\Psi$.*
- (ii) *If ω is dominated by v , then the sets*

$$\{\hat{C}_\Psi(F, \omega) \mid F \text{ is a } v\text{-cofinite } v\text{-}\omega \text{ bond}\}$$

form a basis of open neighbourhoods of ω in $|G|_\Psi$.

Proof. (i) Given a finite bond $F = E(V_1, V_2)$, the set S of its endvertices in V_2 is also finite, and $G[V_1]$ is connected. Hence if $\omega \in \overline{V_1}$, then $\hat{C}_\Psi(F, \omega)$ equals $\hat{C}(S, \omega) \setminus \Psi^0$, which is a basic open neighbourhood of ω in $|G|_\Psi$. Conversely, we need to find for any finite set $S \subseteq V$, without loss of generality connected in G ,⁷ a finite bond F such that $\hat{C}_\Psi(F, \omega) \subseteq \hat{C}(S, \omega)$. As no vertex dominates ω , there is a finite connected set S' of vertices of $C(S, \omega)$ that separates S from ω in G . (Otherwise we could inductively construct an infinite set of disjoint paths in $C(S, \omega)$ each starting at a vertex adjacent to S and ending on some fixed ray $R \in \omega$; then infinitely many of the starting vertices of these paths would share a neighbour in S , which would dominate ω .) As G is finitely separable, there is a finite set of edges separating S from S' in G . As both S and S' are connected, choosing this set minimal ensures that it is a bond. This bond F satisfies $\hat{C}_\Psi(F, \omega) \subseteq \hat{C}(S, \omega)$.

(ii) Although F is infinite now, the set S of its endvertices in V_2 is finite. Hence $\hat{C}_\Psi(F, \omega)$ is a basic open neighbourhood of ω in $|G|_\Psi$, as in the proof of (i). Conversely, let a finite set $S \subseteq V$ be given; we shall find a v -cofinite v - ω bond F such that $\hat{C}_\Psi(F, \omega) \subseteq \hat{C}(S, \omega)$. The sets $\hat{C}(T, \omega)$ such that $v \in T$ and both $T - v$ and T are connected in G still form a neighbourhood basis for ω in $|G|$, so we may assume that S has these properties. As in the proof of (i), there is a finite connected set S' of vertices in $C(S, \omega)$ that separates $S - v$

⁷The sets $\hat{C}(S, \omega)$ with S connected in G also form a neighbourhood basis of ω in $|G|$, since every finite set S of vertices extends to a finite connected set.

from ω in $G - v$, because ω is not dominated in $G - v$. As $G - v$ is finitely separable, there is a finite bond $F = E(V_1, V_2)$ of $G - v$ that separates $S - v$ from S' , with $S - v \subseteq V_2$ say. Then $F' := E(V_1, V_2 \cup \{v\})$ is a v -cofinite v - ω bond in G with $\hat{C}_\Psi(F', \omega) \subseteq \hat{C}(S, \omega)$, as before. \square

Let us call the open neighbourhoods $\hat{C}_\Psi(F, \omega)$ from Lemma 9 the *standard neighbourhoods* in $|G|_\Psi$ of the ends $\omega \in \Psi$. For points of $|G|_\Psi$ other than ends, let their *standard neighbourhoods* be their basic open neighbourhoods defined in Section 2.

Trivially, standard neighbourhoods of vertices and inner points of edges are closed under \sim . Our next lemma says that standard neighbourhoods of ends are nearly closed under \sim , in that only the end itself may be equivalent to points outside: to a vertex dominating it, and to other ends dominated by that vertex.

Lemma 10. *If $\hat{C} = \hat{C}_\Psi(F, \omega)$ is a standard neighbourhood of $\omega \in \Psi$ in $|G|_\Psi$, then $[x]_\Psi \subseteq \hat{C}$ for every $x \in \hat{C} \setminus [\omega]_\Psi$.*

Proof. Let S be the finite set of vertices not in \hat{C} that are incident with an edge in F . Suppose, for a contradiction, that there are points $x \sim y$ in $|G|_\Psi$ such that $x \in \hat{C} \setminus [\omega]$ but $y \notin \hat{C} \setminus [\omega]$. Since the unique vertex in the \sim_Ψ -class of x and y lies either in $\hat{C} \setminus [\omega]$ or not, we may assume that either x or y is that vertex.

Suppose x is the vertex; then y is an end. Let R be a ray of y that avoids S . Then the finite set $S \subseteq V \setminus \{x\}$ separates x from R , a contradiction.

Suppose y is the vertex. If $y \notin S$ we argue as before. Suppose that $y \in S$. Note that y does not dominate ω , since $y \sim x \not\sim \omega$. But now the vertex $v \in S$ that dominates ω , if it exists, and the finitely many neighbours of $S \setminus \{v\}$ in \hat{C} together separate y from every ray in x that avoids this finite set, a contradiction. \square

Let us extend the notion of standard neighbourhoods from $|G|_\Psi$ to \tilde{G}_Ψ . Call a neighbourhood of a point $[x]_\Psi$ of \tilde{G}_Ψ a *standard neighbourhood* if its inverse image under π_Ψ is a union $\bigcup_{y \in [x]_\Psi} U_y$ of standard neighbourhoods U_y in $|G|_\Psi$ of the points $y \in [x]_\Psi$. Neighbourhoods in subspaces of \tilde{G}_Ψ that are induced by such standard neighbourhoods of \tilde{G}_Ψ will likewise be called *standard*. All standard neighbourhoods in \tilde{G}_Ψ and its subspaces are open, by definition of the identification and the subspace topology.

Lemma 11. *For every point $[x]_\Psi \in \tilde{G}_\Psi$ its standard neighbourhoods form a basis of open neighbourhoods in \tilde{G}_Ψ .*

Proof. Given any open neighbourhood N of $[x]_\Psi$ in \tilde{G}_Ψ , its inverse image W under π_Ψ is open in $|G|_\Psi$ and contains every $y \in [x]_\Psi$. By Lemma 9, we can find for each of these y a standard neighbourhood $U_y \subseteq W$ of y in $|G|_\Psi$. By Lemma 10, their union $U = \bigcup_y U_y$ is closed in $|G|_\Psi$ under \sim , so $U = \pi_\Psi^{-1}(\pi_\Psi(U))$. Since U is open in $|G|_\Psi$, this means that $\pi_\Psi(U) \subseteq N$ is an open neighbourhood of $[x]_\Psi$ in \tilde{G}_Ψ . \square

Our next topic is to compare circles in \tilde{G}_Ψ with circles in \tilde{G} . We have already seen that circles in \tilde{G}_Ψ define circles in \tilde{G} , by composition with σ_Ψ . The converse will generally fail: the inverse of σ_Ψ (where it is defined) need not be continuous, so a circle in \tilde{G} need not induce a circle in \tilde{G}_Ψ even if its points all lie in the image of σ_Ψ . This is illustrated by the following example.

Example 3. Consider the graph of Figure 2 with $\Psi = \{\psi\}$. The closure of the double ray D in \tilde{G} is a circle there, since in \tilde{G} the ends ω and ψ are identified. This circle lies in the image of σ_Ψ , but σ_Ψ^{-1} restricted to it fails to be continuous at the point $\{v, \omega, \psi\}$, which σ_Ψ^{-1} maps to the point $\{v, \psi\}$ of \tilde{G}_Ψ .




Figure 3: A circle in \tilde{G} through $p = \{v, \omega, \psi\}$ which defines for $\Psi = \{\psi\}$ a circle in \tilde{G}_Ψ through $\{v, \psi\}$.

However, the map σ_Ψ^{-1} in this example is continuous on the circle in \tilde{G} shown in Figure 3, which ‘does not use’ the end $\omega \in \Psi^{\mathcal{G}}$ when it passes through the point $\{v, \omega, \psi\}$. The fact that circles in \tilde{G} do induce circles in \tilde{G}_Ψ in such cases will be crucial to our proof of Theorem 2:

Lemma 12.

- (i) Let $\rho: S^1 \rightarrow \tilde{G}_\Psi$ be a circle, with image C say, and let D be the set of all inner points of edges on C . Then every end in the $|G|$ -closure⁸ of D lies in Ψ .
- (ii) Let $\varphi: S^1 \rightarrow \tilde{G}$ be a circle, with image C say, and let D be the set of all inner points of edges on C . If every end in the $|G|$ -closure of D lies in Ψ , then the composition $\sigma_\Psi^{-1} \circ \varphi: S^1 \rightarrow \tilde{G}_\Psi$ is well defined and a circle in \tilde{G}_Ψ .

Proof. (i) Consider an end ω in the $|G|$ -closure of D . Since $|G|$ is first-countable by Lemma 4 (unlike \tilde{G} ; see Example 1), there is a sequence $(x_i)_{i \in \mathbb{N}}$ of points in D , chosen from distinct edges, that converges to ω in $|G|$. Suppose $\omega \in \Psi^{\mathcal{G}}$. We show that the x_i have no accumulation point on C , indeed in all of \tilde{G}_Ψ ; this will contradict the fact that C , being a circle, is compact and contains all the x_i .

Consider a point $p \in \tilde{G}_\Psi$. If $p = \{x_n\}$ for some n , then p is not an accumulation point of the x_i , since these lie on distinct edges. Otherwise none of the x_i lies in p .

Let us show that every $y \in p$ has an open neighbourhood O_y in $|G|$ that does not contain any of the x_i . This is clear if y is an inner edge point or a vertex, even a vertex dominating ω , since C contains no more than two edges at any vertex of G and the x_i lie in distinct edges on C . If y is an end, then our assumption of $y \neq \omega$ implies that there is a finite set S of vertices such that $C(S, y) \neq C(S, \omega)$.

⁸We shall freely consider D as a subset of either \tilde{G}_Ψ or $|G|$, and similarly in (ii).

Then all but finitely many x_i lie in $\hat{C}(S, \omega)$, and hence outside $\hat{C}(S, y)$. Adding the endvertices of the finitely many x_i that lie in $\hat{C}(S, y)$ to S , to form S' say, we obtain the desired open neighbourhood $O_y = \hat{C}(S', y)$ of y containing no x_i .

By Lemma 9, each O_y contains a standard neighbourhood U_y of y in $|G|_\Psi$. The union of these is a standard neighbourhood of p in \tilde{G}_Ψ that contains none of the x_i . Hence p is not an accumulation point of the x_i in \tilde{G}_Ψ .

(ii) Assume that every end in the $|G|$ -closure of D lies in Ψ . To show that $\sigma_\Psi^{-1} \circ \varphi$ is well defined, let us prove that $\text{im } \varphi \subseteq \text{im } \sigma_\Psi$. The only points of \tilde{G} not in the image of σ_Ψ are singleton \sim -classes of $|G|$ consisting of an undominated end $\omega \notin \Psi$. By assumption and Lemma 9, such an end ω has a standard neighbourhood in $|G| = |G|_\Omega$ disjoint from D , which π maps to a standard neighbourhood of $\{\omega\}$ in \tilde{G} disjoint from D . So $\{\omega\}$ is not in the \tilde{G} -closure of D . But by [10, Cor. 4.4] that closure is the entire circle C , giving $\{\omega\} \notin \text{im } \varphi$. This completes the proof of $\text{im } \varphi \subseteq \text{im } \sigma_\Psi$. As σ_Ψ is injective, it follows that $\sigma_\Psi^{-1} \circ \varphi$ is well defined.

To show that σ_Ψ^{-1} is continuous on C , let a point $p \in C$ be given. Since p lies in $\text{im } \varphi \subseteq \text{im } \sigma_\Psi$, it is represented by a point x in $G \cup \Psi$; then $p = [x]$ and $\sigma_\Psi^{-1}(p) = [x]_\Psi$. By Lemma 11, it suffices to find for every standard neighbourhood u of $[x]_\Psi$ in $\text{im}(\sigma_\Psi^{-1} \upharpoonright C)$ a neighbourhood w of $[x]$ in C such that $\sigma_\Psi^{-1}(w) \subseteq u$.

By definition, u is the intersection with $\text{im}(\sigma_\Psi^{-1} \upharpoonright C)$ of a set $U \subseteq \tilde{G}_\Psi$ whose inverse image under π_Ψ is a union

$$\pi_\Psi^{-1}(U) = \bigcup_{y \in [x]_\Psi} U_y$$

of standard neighbourhoods U_y in $|G|_\Psi$ of the points $y \in [x]_\Psi$. Our aim is to find a similar set W to define w : a set $W \subseteq \tilde{G}$ such that for $w := W \cap C$ we have $\sigma_\Psi^{-1}(w) \subseteq u$, and such that

$$\pi^{-1}(W) = \bigcup_{y \in [x]} W_y \tag{2}$$

where each W_y is a standard neighbourhood of y in $|G|$.

Let us define these W_y , one for every $y \in [x]$. If $y \in G$, then $y \in [x]_\Psi$. Hence U_y is defined, and it is a standard neighbourhood of y also in $|G|$; we let $W_y := U_y$. If $y \in \Psi$, then again $y \in [x]_\Psi$, and U_y (exists and) has the form $\hat{C}_\Psi(F, y)$. We let $W_y := \hat{C}_\Omega(F, y)$ be its closure in $|G|$; this is a standard neighbourhood of y in $|G|$. Finally, if $y \in \Psi^{\mathfrak{G}}$, then $y \notin [x]_\Psi$ and U_y is undefined. We then let W_y be a standard neighbourhood of y in $|G|$ that is disjoint from D ; this exists by assumption and Lemma 9. Let us call these last W_y *new*.

By Lemma 10, all these W_y are closed under \sim in $|G| \setminus [y]$. Hence $\bigcup_{y \in [x]} W_y$ is closed under \sim in $|G|$. Its π -image W therefore satisfies (2) and is a standard neighbourhood of $[x]$ in \tilde{G} . Hence, $w := W \cap C$ is a neighbourhood of $[x]$ in C .

It remains to show that σ_Ψ^{-1} maps every point $q \in w$ to u . This is clear for $q = p = [x]$, so assume that $q \neq [x]$. By construction of W and Lemma 10,

the set q lies entirely inside one of the W_y . Let us show that no such W_y can be new. Since q is a point in $w \subseteq C$, in which D is dense [10], there is no neighbourhood of q in \tilde{G} that is disjoint from D . But then q has an element z all whose $|G|$ -neighbourhoods meet D . (If not, we could pick for every element of q a standard $|G|$ -neighbourhood disjoint from D ; then the union of all these would project under π to a standard neighbourhood of q in \tilde{G} that avoids D .) As W_y is a $|G|$ -neighbourhood of $z \in q \subseteq W_y$, it thus cannot be new.

We thus have $q \subseteq W_y$ where W_y is the $|G|$ -closure of U_y for some $y \in [x]_\Psi$ (or equal to U_y). In particular, $W_y \setminus U_y \subseteq \Psi^{\mathbf{G}}$. As q lies in C , in which D is dense, we cannot have $q = \{\omega\}$ with $\omega \in \Psi^{\mathbf{G}}$ (as earlier). So either $q = \{\psi\}$ with $\psi \in \Psi$, or q contains a vertex. In either case, $q \cap U_y \neq \emptyset$, which implies that $\sigma_\Psi^{-1}(q) \in U$. As $q \in C$, this implies $\sigma_\Psi^{-1}(q) \in u$, as desired. \square

Lemma 13. *Arc-components of standard subspaces of \tilde{G}_Ψ are closed.*⁹

Proof. Let X be an arc-component of a standard subspace of \tilde{G}_Ψ . If X is not closed, there is a point q in $\tilde{G}_\Psi \setminus X$ such that every (standard) neighbourhood of q meets X . As in the proof of Lemma 12, this implies that q has a representative $y \in |G|_\Psi$ such that every standard neighbourhood U_y of y in $|G|_\Psi$ meets $\pi_\Psi^{-1}(X)$, say in a point $x = x(U_y)$. Clearly, y is an end. Since $x \not\sim y$, we even have $[x]_\Psi \subseteq U_y$ by Lemma 10. Let $U_0 \supseteq U_1 \supseteq \dots$ be a neighbourhood basis for y consisting of such standard neighbourhoods U_y , and let $x_i := x(U_i)$ and $z_i := [x_i]_\Psi$ for all i . Then these x_i converge to y in $|G|_\Psi$, while $(z_i)_{i \in \mathbb{N}}$ is a sequence of points in X that converges in \tilde{G}_Ψ to $q = [y]_\Psi$.

For every $i \in \mathbb{N} \setminus \{0\}$ let A'_i be a z_i - z_0 arc in X . Define subarcs A_i of the A'_i recursively, choosing as A_i the initial segment of A'_i from its starting point z_i to its first point a_i in $\bigcup_{j < i} A_j$, where $A_0 := \{z_0\}$. (The point a_i exists by the continuity of A'_i , since $\bigcup_{j < i} A_j$ is closed, being a compact subspace of the Hausdorff space \tilde{G}_Ψ .)

Define an auxiliary graph H with vertex set $\{A_i \mid i \in \mathbb{N}\}$ and edges $A_i A_j$ whenever j is the smallest index less than i such that $A_i \cap A_j \neq \emptyset$. Suppose first that H has a vertex A_j of infinite degree. Since the arc A_j is compact, it has a point p every neighbourhood of which meets infinitely many A_i . By Lemma 7, there is a finite set F of edges such that in \tilde{G}_Ψ the points p and q have disjoint open neighbourhoods O_p and O_q partitioning $\tilde{G}_\Psi \setminus \dot{F}$. Then for infinitely many i we have both $A_i \cap O_p \neq \emptyset$ and $z_i \in O_q$. For all these i the arc A_i , being connected, must have an edge in the finite set F , a contradiction.

So H is locally finite. By König's infinity lemma, H contains a ray $A_{i_1} A_{i_2} \dots$ such that $i_j < i_k$ whenever $j < k$. We claim that $A := A_{i_1} a_{i_2} A_{i_2} a_{i_3} \dots q$ is an arc in \tilde{G}_Ψ ; this will contradict our assumption that A_{i_1} lies in the arc-component X while q does not. We only have to show that A is continuous at q . Since every neighbourhood of q in \tilde{G}_Ψ contains the π_Ψ -image of one of our standard

⁹This refers to either the subspace or to the entire space \tilde{G}_Ψ ; the two are equivalent, since standard subspaces of \tilde{G}_Ψ are themselves closed in \tilde{G}_Ψ .

neighbourhoods U_n of y , it suffices to show that for every such U_n we have $A_i \subseteq \pi_\Psi(U_n)$ for all but finitely many i .

Since U_n is a standard neighbourhood of y , there exists a set F of edges such that $U_n = \hat{C}_\Psi(F, y)$ and F is either finite or v -cofinite with $v \sim y$. Since none of the A_i contains such a vertex v , and only finitely many of them meet \hat{F} , all but finitely many A_i lie in $(\tilde{G}_\Psi \setminus \{q\}) \setminus \hat{F}$ and have their starting point $z_i = [x_i]_\Psi$ in $\pi_\Psi(U_n)$, by the choice of U_n . To complete our proof, we shall show that $\pi_\Psi(U_n \setminus q) \setminus \hat{F}$ and its complement in $(\tilde{G}_\Psi \setminus \{q\}) \setminus \hat{F}$ are two open subsets of $(\tilde{G}_\Psi \setminus \{q\}) \setminus \hat{F}$ partitioning it: then none of those cofinitely many A_i can meet both, so they will all lie entirely in $\pi_\Psi(U_n)$.

Since U_n is a standard neighbourhood of $y \in q$, the set $U_n \setminus q$ is open in $|G|_\Psi \setminus q$ and closed under \sim , so $\pi_\Psi(U_n \setminus q)$ is open in $\tilde{G}_\Psi \setminus \{q\}$ and $\pi_\Psi(U_n \setminus q) \setminus \hat{F}$ is open in $(\tilde{G}_\Psi \setminus \{q\}) \setminus \hat{F}$. Its complement in $(\tilde{G}_\Psi \setminus \{q\}) \setminus \hat{F}$ is open, because it is the π_Ψ -image of the (\sim -closed) union of the finite set S of vertices that are incident with edges in F but are not in U_n , the edges incident with them that are not in \hat{F} , and the $|G|_\Psi$ -closures of the components of $G - S$ not contained in U_n . The two open sets partition all of $(\tilde{G}_\Psi \setminus \{q\}) \setminus \hat{F}$, because U_n is itself the $|G|_\Psi$ -closure of a component of $G - S$ together with the edges between S and that component (which all lie in F). \square

4 Proof of Theorem 2

We can now apply the lemmas from Section 3 to prove Theorem 2. One of these lemmas, Lemma 12, also implies a characterization of duality in terms of circuits and bonds. In order to include this in the statement of the theorem, we need another definition. If Ψ is a set of ends of a graph G , a bond D of G is a Ψ -bond if every end of G that lies in the closure of $\bigcup D$ in $|G|$ is in Ψ .

Theorem 14. *Let $G = (V, E, \Omega)$ and $G^* = (V^*, E, \Omega)$ be finitely separable 2-connected graphs with the same edge set E and the same end set Ω . Then the following assertions are equivalent:*

- (i) G and G^* are duals of each other.
- (ii) For all $\Psi \subseteq \Omega$ and $F \subseteq E$ the following holds: F is the edge set of a Ψ -tree in G if and only if F^G is the edge set of a Ψ^G -tree in G^* .
- (iii) There exists a set $\Psi \subseteq \Omega$ such that for every $F \subseteq E$ the following holds: F is the edge set of a Ψ -tree in G if and only if F^G is the edge set of a Ψ^G -tree in G^* .
- (iv) For all $\Psi \subseteq \Omega$, the Ψ -circuits of G are precisely the Ψ -bonds of G^* .
- (v) There exists a set $\Psi \subseteq \Omega$ such that the Ψ -circuits of G are precisely the Ψ -bonds of G^* .

Remark. The fact that (i)–(iii) are symmetrical in G and G^* , while (iv) and (v) are not, is immaterial and only serves to avoid clutter: as noted before, it was proved in [4, Theorem 3.4] that if G^* is a dual of G then G is a dual of G^* .

We shall prove the implications (i)→(iv)→(v)→(i) first, and then the implications (i)→(ii)→(iii)→(i). The two proofs can be read independently.

(i)→(iv) Assume (i), and let $\Psi \subseteq \Omega$ and $D \subseteq E$ be given for a proof of (iv). If D is a Ψ -circuit of G , for the circle $\rho: S^1 \rightarrow \tilde{G}_\Psi$ say, it is also a circuit of G with circle $\sigma_\Psi \circ \rho: S^1 \rightarrow \tilde{G}$. By (i), then, D is a bond of G^* . By Lemma 12 (i), every end in the closure of $\bigcup D$ lies in Ψ , so D is in fact a Ψ -bond of G^* .

Conversely, if D is a Ψ -bond of G^* , then D is a circuit of G by (i), with circle $\varphi: S^1 \rightarrow \tilde{G}$ say, and by Lemma 12 (ii) the composition $\sigma_\Psi^{-1} \circ \varphi$ is a circle in \tilde{G}_Ψ . The edges it contains are precisely those in D , so D is a Ψ -circuit.

(iv)→(v) Using the empty set for Ψ in (iv) immediately yields (v).

(v)→(i) By [4, Lemma 4.7 (i)], G^* is dual to G as soon as the *finite* circuits of G are precisely the finite bonds of G^* . This is immediate from (v).

Let us now prove the implications (i)→(ii)→(iii)→(i). When we consider edges in E topologically, we take them to include their endvertices in \tilde{G}_Ψ or in $\tilde{G}_{\Psi^c}^*$, depending on the context. Thus, in (ii) and (iii), $\bigcup F$ will be a subspace of \tilde{G}_Ψ while $\bigcup F^c$ will be a subspace of $\tilde{G}_{\Psi^c}^*$.

(i)→(ii) We first show that (i) implies the analogue of (ii) with ordinary topological connectedness, rather than the arc-connectedness required of a Ψ -tree:

(\star) For all $F \subseteq E$ and $\Psi \subseteq \Omega$: F is the edge set of a connected spanning Ψ -forest of G if and only if F^c is the edge set of a connected spanning Ψ^c -forest of G^* .

For our proof of (\star) from (i), let $F \subseteq E$ and $\Psi \subseteq \Omega$ be given, and assume that F is the edge set of a connected spanning Ψ -forest T of G . Let X be the closure in $\tilde{G}_{\Psi^c}^*$ of $V(\tilde{G}_{\Psi^c}^*) \cup \bigcup F^c$. We shall prove that X is a connected subspace of $\tilde{G}_{\Psi^c}^*$ that contains no circle. Then X will be a standard subspace, and it is spanning by definition. Roughly, the idea is that X should be connected because T is acirclic, and acirclic because T is connected.

Let us show first that X contains no circle. Suppose there is a circle $\varphi: S^1 \rightarrow X$, with circuit $D \subseteq F^c$ say. By Lemma 12 (i) applied to G^* and Ψ^c , every end in the $|G^*|$ -closure of $\bigcup D$ lies in Ψ^c . But the ends in the $|G^*|$ -closure of $\bigcup D$ are precisely those in its $|G|$ -closure, by (i). Hence we obtain:

The $|G|$ -closure of $\bigcup D$ contains no end from Ψ . (3)

Since D is also the circuit of the circle $\sigma_{\Psi^c} \circ \varphi: S^1 \rightarrow \tilde{G}^*$, assumption (i) implies that D is a bond in G ; let $\{V_1, V_2\}$ be the corresponding partition of V . Let us show the following:

Every point $p \in \tilde{G}_\Psi$ has a standard neighbourhood N such that $\pi_\Psi^{-1}(N)$ contains vertices from at most one of the sets V_1 and V_2 . (4)

Suppose $p \in \tilde{G}_\Psi$ has no such neighbourhood. Then p has a representative x all whose standard neighbourhoods in $|G|_\Psi$ meet V_1 , and a representative y all whose standard neighbourhoods in $|G|_\Psi$ meet V_2 .

If $x = y$, the point $x = y =: \psi$ is an end in Ψ . Then every standard neighbourhood of ψ in $|G|_\Psi$ contains a graph-theoretical path from V_1 to V_2 , and hence an edge from D , because the subgraphs of G underlying standard neighbourhoods in $|G|_\Psi$ are connected and meet both V_1 and V_2 . This contradicts (3).

So $x \neq y$. In particular, p is nontrivial, so it contains a vertex v , say in V_1 . Then $v \neq y$, so $y =: \psi \in \Psi$. Pick a ray $R \in \psi$. Replacing R with a tail of R if necessary, we may assume by (3) that R has no edge in D . If all the vertices of R lie in V_1 , then every standard neighbourhood of $y = \psi$ meets both V_1 and V_2 , which contradicts (3) as in the case of $x = y$. So $R \subseteq G[V_2]$. Let us show that every standard neighbourhood $\hat{C}_\Psi(F', \psi)$ of ψ contains the inner points of an edge from D , once more contrary to (3).

By Lemma 9 (ii), F' is v -cofinite. Since $v \sim \psi$, there are infinitely many v - R paths P_0, P_1, \dots in G that meet pairwise only in v . Since D separates v from R , each P_i contains an edge $e_i \in D$. Only finitely many of the P_i contain one of the finitely many edges from F' that are not incident with v . All the other P_i have all their points other than v in $\hat{C}_\Psi(F', \psi)$, including the inner points of e_i . This completes the proof of (4).

For every point $p \in \tilde{G}_\Psi$ pick a standard neighbourhood N_p as in (4). Let O_1 be the union of all edges with an endvertex in V_1 and all those N_p such that $\pi_\Psi^{-1}(N_p)$ meets V_1 , and define O_2 likewise for V_2 . Then O_1, O_2 are two open subsets of \tilde{G}_Ψ covering it, and it is easy to check that $O_1 \cap O_2 \subseteq \mathring{D}$. So no connected subspace of $\tilde{G}_\Psi \setminus \mathring{D}$ contains vertices from V_1 as well as from V_2 . But our connected spanning Ψ -forest T is such a subspace, since its edges lie in $F \subseteq E \setminus D$. This contradiction completes the proof that X contains no circle.

For the proof of (\star) it remains to show that X is connected. If not, there are open sets O_1, O_2 in $\tilde{G}_{\Psi^c}^*$ that each meet X and together cover it, but intersect only outside X . It is easy to check that, since X contains all the vertices of $\tilde{G}_{\Psi^c}^*$, both O_1 and O_2 contain such a vertex but they have none in common. For $i = 1, 2$, let V_i^* be the set of vertices of G^* representing a vertex of $\tilde{G}_{\Psi^c}^*$ in O_i . Let C be a bond contained in the cut $E(V_1^*, V_2^*)$. Note that the edges e of this bond all lie in F : as e is connected but contained in neither O_i , it cannot lie in $O_1 \cup O_2 = X$. As F is the edge set of a Ψ -forest, $C \subseteq F$ cannot be a Ψ -circuit of G . By (i), however, C is a circuit of G , because it is a bond of G^* . By Lemma 12 (ii), therefore, there is an end $\omega \in \Psi^c$ in the $|G|$ -closure of \hat{C} ; then ω also lies in the $|G^*|$ -closure of \hat{C} .

Let us show that every standard neighbourhood W of $[\omega]_{\Psi^c}$ in $\tilde{G}_{\Psi^c}^*$ contains an edge from C , including its endvertices in $\tilde{G}_{\Psi^c}^*$. By definition, W is the image under π_{Ψ^c} of a subset of $|G^*|_{\Psi^c}$ that contains a standard neighbourhood U of ω in $|G^*|_{\Psi^c}$. Since ω lies in the $|G^*|$ -closure of \hat{C} , this U either contains infinitely many edges from C together with their endvertices in G^* , or it contains the inner points and one endvertex (in G^*) of infinitely many edges from C whose other endvertex dominates ω in G^* . In both cases, these edges $e \in C$ and their endvertices in $\tilde{G}_{\Psi^c}^*$ lie in W .

So every standard neighbourhood of $[\omega]_{\Psi^c}$ in $\tilde{G}_{\Psi^c}^*$ contains an edge from C ,

including its endvertices in $\widetilde{G}_{\Psi^{\mathcal{C}}}^*$. In particular, it meets X in both O_1 and O_2 , where this edge has its endvertices. So every neighbourhood of $[\omega]_{\Psi^{\mathcal{C}}}$ in X meets both O_1 and O_2 . This contradicts the fact that the O_i induce disjoint open subsets of X of which only one contains the point $[\omega]_{\Psi^{\mathcal{C}}}$. This completes the proof of (\star) .

It remains to derive the original statement of (ii) from (\star) . Suppose (ii) fails, say because there is a Ψ -tree T of G , with edge set F say, such that $F^{\mathcal{C}}$ is not the edge set of a $\Psi^{\mathcal{C}}$ -tree of G^* . By (\star) we know that $F^{\mathcal{C}}$ is the edge set of a connected spanning $\Psi^{\mathcal{C}}$ -forest X in G^* , which we now want to show is even arc-connected. Suppose it is not. Since the arc-components of X are closed (Lemma 13), no arc-component of X contains all its vertices. Vertices in different arc-components are joined by a finite path in G^* , which contains an edge e whose endvertices lie in different arc-components of X . Then $X \cup e$ still contains no circle, so $F^{\mathcal{C}} \cup \{e\}$ too is the edge set of a connected spanning $\Psi^{\mathcal{C}}$ -forest of G^* . Thus, by (\star) , $F \setminus \{e\}$ is the edge set of a connected spanning Ψ -forest of G . This can only be $T \setminus \{e\}$, so $T \setminus \{e\}$ has precisely two arc-components D_1 and D_2 but is still connected. Then D_1 and D_2 cannot both be open, or equivalently, cannot both be closed. This contradicts Lemma 13.

(ii) \rightarrow (iii) Using the empty set for Ψ in (ii) immediately yields (iii).

(iii) \rightarrow (i) By [4, Lemma 4.7 (i)] it suffices to show that G^* is a finitary dual of G , i.e., that the finite circuits of G are precisely the finite bonds of G^* . By Lemma 6 (ii), a finite set F of edges is a bond of G^* if and only if it meets the edge set of every $\Psi^{\mathcal{C}}$ -tree of G^* and is minimal with this property. By (iii), this is the case if and only if F is not contained in the edge set of any Ψ -tree of G , and is minimal with this property. By Lemma 6 (i), this is the case if and only if F is a circuit of G . \square

5 Outlook: geometric duality and matroids

The aim of this paper was to generalize to infinite graphs the equivalence, long known for finite graphs, between the cycle-bond duality that comes from topology and the more algebraic spanning tree duality that comes from matroids. We have shown that this works very generally with the Ψ -trees of G and the $\Psi^{\mathcal{C}}$ -trees of G^* , for all subsets Ψ of the ends of G and G^* . But although the notion of Ψ -trees was inspired by matroid duality, we have not addressed the question of whether the Ψ -trees of a graph, even a locally finite one, do indeed always come from a matroid.

Bowler and Carmesin have studied this question in some depth [1, 2]. They came up with a surprising answer: there exist examples of locally finite graphs (G, E, Ω) and sets $\Psi \subseteq \Omega$ for which the Ψ -trees of G do *not* form the bases of a matroid. But such examples are necessarily esoteric [1].¹⁰ When the Ψ -trees of G do define a matroid, i.e. in most cases, they also have the sort of properties one would expect of spanning trees: their edge sets are maximal without con-

¹⁰The set Ψ must fail to be Borel in Ω . Naturally definable subsets $\Psi \subseteq \Omega$ are usually Borel.

taining a Ψ -circuit, and minimal with the property of meeting every $\Psi^{\mathcal{G}}$ -bond. Similarly, Lemma 6 then extends to infinite edge sets F with Ψ -circuits and $\Psi^{\mathcal{G}}$ -bonds.

Conversely, Bowler and Carmesin [2] showed that all (tame) matroids whose circuits are edge sets of circles in $|G|$ and whose cocircuits are bonds of G are Ψ -matroids for some Ψ . Thus, although there are many ways of specifying collections of circles in $|G|$, whenever these form a matroid they can be defined simply by specifying which ends of G those circles are allowed to use.

The Ψ -trees we introduced merely for the purpose of graph duality, therefore, play a fundamental role also for infinite graphic matroids in general – as long as ‘graphic’ means that the matroid’s circuits are the edge sets of circles in $|G|$ for some graph G . (There is also the *algebraic cycle matroid* of a graph, whose infinite circuits are the edge sets of its double rays; this is well understood [5]. And there are matroids that are ‘graphic’ merely in the sense that their finite minors are graphic. These need not have the edge sets of circles of a graph as their circuits, but their circuits can be represented by circles in more general topological spaces [3].)

The sets Ψ for which the Ψ -trees of a locally finite graph G do not give a matroid have an interesting property: there are subsets of $|G|_{\Psi}$ that are closed (in $|G|_{\Psi}$, not in $|G|$) and connected but not arc-connected.¹¹ Might it be the case that, for such Ψ , relaxing the requirement of arc-connectedness in the definition of Ψ -trees to topological connectedness can mend their failure to give a matroid?

Given $G = (V, E, \Omega)$ and $\Psi \subseteq \Omega$, call a set $F \subseteq E(G)$ a *topological Ψ -tree* in G if every $v \in V$ is incident with an edge in F , the closure of $\bigcup F$ in $|G|_{\Psi}$ is topologically connected, and F is \subseteq -minimal with this last property.¹²

Problem. *Do the topological Ψ -trees in a locally finite graph always form the bases of a matroid?*

Finally, one may ask whether pairs of dual finitely separable graphs have geometric representations as known from finite graphs:

Embedding Conjecture. Let $G = (V, E, \Omega)$ and $G^* = (V^*, E^*, \Omega^*)$ be dual 2-connected finitely separable graphs, with compatible bijections $E \rightarrow E^*$ and $\Omega \rightarrow \Omega^*$, let $\Psi \subseteq \Omega$, and put $\Psi^* := \{\omega^* \mid \omega \in \Psi^{\mathcal{G}}\}$. Then there exist embeddings of \tilde{G}_{Ψ} and $\tilde{G}_{\Psi^*}^*$ in the plane that meet exactly in the midpoints of corresponding edges e, e^* .

Acknowledgement

The ideas that led to the formulation of Theorem 2 were developed jointly with Henning Bruhn. We benefited greatly from his insights at this stage. We also benefited from the comments of a fantastically thorough referee, which prompted us to explain a number of technical arguments better or in more detail.

¹¹This follows from [1, Thms 4.2 & 4.4] and the subsequent discussion. Compare also [11].

¹²For all Borel sets Ψ , the topological Ψ -trees in G are precisely the edge sets of its Ψ -trees.

References

- [1] N. Bowler and J. Carmesin, *Infinite matroids and determinacy of games*, arXiv:1301.5980, 2013.
- [2] ———, *The ubiquity of Ψ -matroids*, arXiv:1304.6973, 2013.
- [3] N. Bowler, J. Carmesin, and R. Christian, *Infinite graphic matroids, Part I*, arXiv:1309.3735, 2013.
- [4] H. Bruhn and R. Diestel, *Duality in infinite graphs*, *Comb., Probab. Comput.* **15** (2006), 75–90.
- [5] ———, *Infinite matroids in graphs*, *Discrete Math.* **311** (2011), 1461–1471, Special volume on *Infinite Graphs: Introductions, Connections, Surveys* (R. Diestel, G. Hahn & B. Mohar, eds).
- [6] Henning Bruhn and Maya Stein, *Duality of ends*, *Comb., Probab. Comput.* **19** (2010), 47–60.
- [7] R. Diestel, *End spaces and spanning trees*, *J. Combin. Theory (Series B)* **96** (2006), 846–854.
- [8] ———, *Graph Theory*, 4th ed., Springer, 2010.
- [9] R. Diestel, *Locally finite graphs with ends: a topological approach*, *Discrete Math.* **310–312** (2010–11), 2750–2765 (310); 1423–1447 (311); 21–29 (312), arXiv:0912.4213.
- [10] R. Diestel and D. Kühn, *Topological paths, cycles and spanning trees in infinite graphs*, *Europ. J. Combinatorics* **25** (2004), 835–862.
- [11] A. Georgakopoulos, *Connected but not path-connected subspaces of infinite graphs*, *Combinatorica* **27** (2007), no. 6, 683–698.
- [12] D.W. Hall and G.L. Spencer, *Elementary topology*, John Wiley, New York 1955.
- [13] C. Thomassen, *Planarity and duality of finite and infinite graphs*, *J. Combin. Theory (Series B)* **29** (1980), 244–271.
- [14] ———, *Duality of infinite graphs*, *J. Combin. Theory (Series B)* **33** (1982), 137–160.