

Duality in infinite graphs

Henning Bruhn Reinhard Diestel

Abstract

The adaption of combinatorial duality to infinite graphs has been hampered by the fact that while cuts (or cocycles) can be infinite, cycles are finite. We show that these obstructions fall away when duality is reinterpreted on the basis of a ‘singular’ approach to graph homology, whose cycles are defined topologically in a space formed by the graph together with its ends and can be infinite. Our approach enables us to complete Thomassen’s results about ‘finitary’ duality for infinite graphs to full duality, including his extensions of Whitney’s theorem.

1 Introduction

The *cycle space* over \mathbb{Z}_2 of a finite graph G is the set of all symmetric differences of its *circuits*, the edge sets of the cycles in G . If G^* is another graph and there exists a bijection $E(G) \rightarrow E(G^*)$ that maps the circuits of G precisely to the minimal non-empty cuts (or *bonds*) of G^* , then G^* is called a *dual* of G . The classical result in this context is Whitney’s theorem:

Theorem 1.1 (Whitney [14]). *A finite graph G has a dual if and only if it is planar.*

For infinite graphs, however, there is an obvious asymmetry between circuits and cuts that gets in the way of duality: while cuts can be infinite, cycles are finite. Indeed, let G be the half-grid shown in solid lines in Figure 1. Geometrically, the dotted graph G^* should be its dual. But then various infinite sets of edges in G , such as the edge sets of its horizontal 2-way infinite paths, should be circuits, because they correspond to bonds of G^* .




Figure 1: A dual with infinite bonds requires infinite circuits

This problem ties in with similar recently observed difficulties about extending other homology aspects of finite graphs to infinite graphs [6]. In all those

cases, the problems could be resolved in a topologically motivated extension of the cycle space, which takes as its basic circuits not only the edge sets of finite cycles but more generally the edge sets of all *circles* – the homeomorphic images of S^1 in the space $|G|$ formed by the graph together with its ends. (When G is locally finite, $|G|$ is known as its *Freudenthal compactification*. In topological terms, the key step is to use a singular-type homology in $|G|$ rather than the simplicial homology of the graph G .) In our example, every 2-way infinite horizontal path D in G forms such a circle in $|G|$ together with the unique end of G , since both its tails converge to this end. The edge set of D would thus be a circuit, as required by its duality to a bond of G^* .

We show that the obstructions to duality in infinite graphs can indeed be overcome in this way. As in those other cases, we build the cycle space not just on the edge sets of the finite cycles of G but on the edge sets of all its circles. This time, however, we take these circles not in $|G|$ itself but in a natural quotient space of $|G|$: whenever an infinite path R is *dominated* by a vertex v (that is, if G contains infinitely many v - R paths that are disjoint except in v) we identify v with the end of R , so that R converges to v rather than to a newly added point at infinity. In Figure 1, every infinite path in the dotted graph converges to the vertex v , so every maximal horizontal path and every maximal vertical path in G^* forms a circle through v .

Our results extend the duality of finite graphs in what appears to be a complete and best-possible way. However, all our results build on previous work of Thomassen [12, 13], who extended finite duality to infinite graphs as far as it will go without considering infinite circuits. Thomassen’s approach was to overcome the disparity between finite cycles and infinite cuts by disregarding the latter. In his terms, G^* qualifies as a dual of G as soon as the edge bijection between the two graphs maps all the (finite) circuits of G to the finite bonds of G^* .

This weaker notion of duality already permits quite a satisfactory extension of Whitney’s theorem to 2-connected graphs. Our stronger notion strengthens this (in one direction), and in addition re-establishes two aspects of finite duality that cannot be achieved for infinite graphs when infinite cuts (and cycles) are disregarded: the uniqueness of duals for 3-connected graphs, and symmetry, the fact that a graph is always a dual of its duals.

Concerning symmetry, note that taking duals may force us out of the class of locally finite graphs: while the graph G in Figure 1 is locally finite, its dual is not. But graphs with vertices of infinite degree do not, in general, have duals. Indeed, Thomassen showed that any infinite graph with a dual (even in his weaker sense) must satisfy the following condition:

No two vertices are joined by infinitely many edge-disjoint paths. (*)

To achieve symmetry, we thus need that the class of graphs satisfying (*) is closed under taking duals. While this is not the case for Thomassen’s notion, it will be true for ours.

Our paper is organised as follows. After providing the required terminology in Section 2 we continue the above discussion in more precise terms in Section 3, which leads up to the statement of our basic duality theorem, Theorem 3.4. We prove this theorem in Section 4. In Section 5 we characterise the graphs that have locally finite duals. In Section 6 we treat duality in terms of spanning trees. In Section 7 we apply our results to colouring-flow duality.

2 Definitions

The basic terminology we use can be found in [7]. However, *graphs* in this paper are the ‘multigraphs’ of [7]: they may have loops and multiple edges. When we contract an edge $e = uv$ then this may create loops (from edges parallel to e) and new parallel edges (if u and v had a common neighbour). Observe that contracting a loop is the same as deleting it. Following Thomassen [12, 13], we require 2-connected graphs to be loopless, and 3-connected graphs to have no parallel edges.¹ A 1-way infinite path is called a *ray*, a 2-way infinite path is a *double ray*, and the subrays of a ray or double ray are their *tails*.

A graph is said to be *planar* if it can be drawn in the plane in such a way that if the images of two edges e and f have a point in common then this point corresponds to a vertex incident with both e and f , and such that all the images of vertices are disjoint. By Kuratowski’s theorem and compactness, a countable graph is planar if and only if it contains neither K_5 nor $K_{3,3}$ as a minor [10].

Let $G = (V, E)$ be a graph, 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 $\Omega = \Omega(G)$. An end ω is said to be *dominated* by a vertex v if for some (equivalently: for every) ray $R \in \omega$ there are infinitely many v – R paths that meet only in v . The ends of G that are not dominated are precisely its topological ends as defined by Freudenthal [11]; see [8] for details.

We will now define two topological spaces. The first of these, denoted as $|G|$, has $V \cup \Omega \cup \bigcup E$ as its point set; we shall call its topology VTOP. The other, denoted as \tilde{G} , will be a quotient space of $|G|$. Its point set can be viewed as $V \cup \Omega' \cup \bigcup E$, where Ω' is the set of undominated ends, and we call its topology ITOP. When G is locally finite, the two spaces will coincide.

Let us start with $|G|$. We begin by viewing G itself (without ends) as the point set of a 1-complex. Then every edge is a copy of the real interval $[0, 1]$, and we give it the corresponding metric and topology. For every vertex v we take as a basis of open neighbourhoods the open stars of radius $1/n$ around v . (That is to say, for every integer $n \geq 1$ we declare as open the set of all points on edges at v that have distance less than $1/n$ from v , in the metric of that edge.)² In order to extend this topology to Ω , we take as a basis of open neighbourhoods of a given end $\omega \in \Omega$ the sets of the form

$$C(S, \omega) \cup \Omega(S, \omega) \cup \overset{\circ}{E}(S, \omega),$$

where $S \subseteq V$ is a finite set of vertices, $C(S, \omega)$ is the unique component of $G - S$ in which every ray from ω has a tail, $\Omega(S, \omega)$ is the set of all ends $\omega' \in \Omega$ whose rays have a tail in $C(S, \omega)$, and $\overset{\circ}{E}(S, \omega)$ is the set of all inner points of edges between S and $C(S, \omega)$.³ We shall freely view G and its subgraphs either as

¹This is motivated by matroid theory. Disallowing loops is also necessary for uniqueness: a graph with loops never has exactly one dual (unless it is itself a loop).

²If G is locally finite, this is the usual identification topology of the 1-complex. Vertices of infinite degree, however, have a countable neighbourhood basis in VTOP, which they do not have in the 1-complex.

³In the early papers on this topic, such as [9], some more basic open sets were allowed: in the place of $\overset{\circ}{E}(S, \omega)$ we could take an arbitrary union of open half-edges from C towards S , one from every S – C edge. When G is locally finite, this yields the same topology. When G has vertices of infinite degree, our topology is slightly sparser but still yields the same cycle space; see the end of this section for more discussion.

abstract graphs or as subspaces of $|G|$. Note that in $|G|$ every ray converges to the end of which it is an element.

Now let \tilde{G} be the quotient space obtained from $|G|$ by identifying each vertex v with all the ends it dominates. Since G (ie., all the graphs we shall consider) will satisfy $(*)$, the equivalence class containing v contains no other vertex. We also denote this class by v , and think of it as the old vertex v to which now the rays dominated by v converge. This quotient space \tilde{G} is easily seen to be Hausdorff (unlike $|G|$, where we cannot find disjoint open sets for an end and a vertex that dominates it), and if G is 2-connected then \tilde{G} is compact [5].

For the definitions that follow we shall formally work in \tilde{G} , but bear in mind that they apply also to $|G|$ when G is locally finite (in which case no identification takes place and $\tilde{G} = |G|$).

A subset of \tilde{G} is a *circle* (respectively, an *arc*) if it is homeomorphic to the unit circle S^1 in the Euclidean plane (respectively, to the real interval $[0, 1]$). For example, every horizontal double ray in the graph G of Figure 1 forms, together with the unique end of G , a circle in $\tilde{G} = |G|$, because its tails converge to this end. Similarly, every vertical ray in the dotted graph G^* that starts at v forms a circle in \tilde{G}^* (but not in $|G^*|$), because in ITOP its end – the unique end of G^* – is identified with its starting vertex v .

Note that a circle C includes every edge of which it contains an inner point, and thus it has a well-defined edge set, called its *circuit*. Conversely, it is not hard to show [9] that $C \cap G$ is dense in C , so every circle is the closure in \tilde{G} of the union of the edges in its circuit, and hence defined uniquely by its circuit. Note that every finite circuit in G is also a circuit in this sense.

Call a family $(D_i)_{i \in I}$ of subsets of E *thin* if no edge appears in infinitely members of the family. Let the *sum* $\sum_{i \in I} D_i$ of this family be the set of all edges that lie in D_i for an odd number of indices i . Then the *cycle space* $\mathcal{C}(\tilde{G})$ of \tilde{G} is the set of all sums of (thin families of) edge sets of circuits, finite or infinite. Symmetric difference as addition makes $\mathcal{C}(\tilde{G})$ into a \mathbb{Z}_2 vector space, which coincides with the usual cycle space of G over \mathbb{Z}_2 when G is finite. We remark that $\mathcal{C}(\tilde{G})$ is closed also under taking infinite thin sums [9], which is not obvious from the definitions. When G is finite or locally finite, we usually write $\mathcal{C}(G)$ instead of $\mathcal{C}(\tilde{G})$.

A set $F \subseteq E$ is a *cut* of G if there is a partition (A, B) of V such that F is the set of all the edges of G with one vertex in A and the other in B . We shall also denote this set by $E(A, B)$. A cut is called a *bond* if it is minimal among the non-empty cuts.

We shall need the following two results as tools in our proofs.

Theorem 2.1. [9] *Let G be a graph satisfying $(*)$. Then every element of $\mathcal{C}(\tilde{G})$ is a disjoint union of circuits.*

Theorem 2.2. [9] *Let G be a graph satisfying $(*)$. Then a set $Z \subseteq E(G)$ is an element of $\mathcal{C}(\tilde{G})$ if and only if Z meets every finite cut in an even number of edges.*

For the conscientious reader we remark that, although the topology for $|G|$ considered in [9] is slightly larger than ours (see the earlier footnote), the above two theorems are nevertheless applicable in our context. This is because the circuits in \tilde{G} coincide for these topologies: as one readily checks, the identity on

\tilde{G} between the two spaces is bicontinuous when restricted to a circle in either space.

3 Duality in infinite graphs

As discussed in the introduction, Thomassen pursued an approach to duality in infinite graphs that is based solely on finite circuits and cuts. While being very successful in some respects, such as conditions for the existence of duals and extensions of Whitney’s theorem, this approach leads to unavoidable problems in others, such as symmetry and the uniqueness of duals. Our aim in this section is to discuss these problems, to indicate why considering infinite circuits and working in ITOP is both, in essence, necessary and sufficient to cure them, and to state our main result, Theorem 3.4.

Consider graphs G and G^* , possibly infinite, and assume that there is a bijection $*$: $E(G) \rightarrow E(G^*)$. Given a set $F \subseteq E(G)$, put $F^* := \{e^* \mid e \in F\}$, and vice versa. (That is, given a subset of $E(G^*)$ denoted by F^* , we write F for the subset $\{e \mid e^* \in F^*\}$ of $E(G)$.) Call G^* a *finitary dual* of G if, for every finite set $F \subseteq E(G)$, the set F is a circuit in G if and only if F^* is a bond in G^* .

Expressed in these terms, Thomassen obtained the following extension of Whitney’s theorem:

Theorem 3.1 (Thomassen [13]). *A 2-connected⁴ graph G has a finitary dual if and only if G is planar and satisfies $(*)$.*

Going back to Figure 1, we see that the dotted graph G^* is a finitary dual of the half-grid G . However, splitting the vertex v into two vertices u and w , and making each of these adjacent to infinitely many neighbours of v in such a way that every neighbour of v is adjacent to exactly one of u and w , we obtain another finitary dual H of G . This violates the intended uniqueness of duals for 3-connected graphs such as G . (Recall that duals of 3-connected finite graphs are unique.)

Moreover, admitting H as a dual violates symmetry, since G is not a finitary dual of H . In fact, H has no finitary dual at all, and it might not even be planar, depending on how we join u and w to the neighbours of v .

Thomassen realised these problems, as is witnessed by the following two theorems.

Theorem 3.2 (Thomassen [12]). *Let G be a 2-connected graph having a finitary dual. Then G has a finitary dual G^* satisfying $(*)$, and every such finitary dual G^* has G as its finitary dual.*

We say that a graph H is a *finitary predual* of G if G is a finitary dual of H .

Theorem 3.3 (Thomassen [13]). *If G has a 3-connected finitary predual then this is its only predual, up to isomorphism.*

⁴We expect that Thomassen’s theorem extends to graphs of smaller connectivity. There is no mention of this in [13], however, and we note that the canonical proof for the forward implication fails: when G^* is a finitary dual of G , then the duality map $*$ need not map the blocks of G to blocks of G^* , so it is not obvious that the blocks of G have finitary duals too. Compare Lemma 4.8 below.

By considering infinite as well as finite circuits, however, we can restore uniqueness. In our example, consider the edge set F of the double ray D in G . In G^* , its dual set F^* (the set of edges incident with v) is a bond. But F^* is not a bond in H , because it contains the edges incident with u (say) as a proper subset. Thus, if F counts as a circuit, then G^* will be a dual of G but H will not, as should be our aim. Taking the circuits of G in $|G| = \tilde{G}$ achieves this.

To restore symmetry, we have to allow vertices of infinite degree. (Note that G^* has one, and we want G to be its dual.) We thus have to decide now whether to work in $|G|$ or in \tilde{G} . That is to say, should we take the circles that define our infinite circuits in the topology ITOP specifically designed for graphs satisfying (*), or in the simpler VTOP?

To answer this question, let us consider the graph G shown in unbroken lines in Figure 2, and let G^* be a hypothetical dual of G (by the definition we are seeking). We want G to be a dual of G^* , and in particular a finitary dual. Thus, G^* will be a finitary predual of G . Now G is certainly a finitary dual of the dotted graph H shown in Figure 2, so H is also a finitary predual of G . Since H is 3-connected, Theorem 3.3 implies that $H = G^*$.




Figure 2: The self-dual graph G

Since the dotted edges at v form a bond of G^* , we thus have to make the edge set F of the double ray D a circuit of G . Now in $|G|$ the set F is not a circuit, because G has two ends and D has a tail in each. In \tilde{G} , however, both ends of G are identified with the vertex v , so the double ray D and the vertex v together do form a circle, making F into a circuit as desired.

We therefore propose the following stronger notion of duality for infinite graphs, in which the duality condition is required of all sets of edges, finite or infinite, and circuits are defined as in \tilde{G} under ITOP.

Definition. Let G be a graph satisfying (*). Let G^* be another graph, with a bijection $*$: $E(G) \rightarrow E(G^*)$. Call G^* a dual of G if the following holds for every set $F \subseteq E(G)$, finite or infinite: F is a circuit in \tilde{G} if and only if F^* is a bond in G^* .

Note that every dual in this sense is also a finitary dual, but not conversely.

Figure 3 shows that infinite circuits can get pretty wild, even in locally finite graphs. The following theorem, which is our main result, can thus deviate more from the corresponding finite situation than it might at first appear.




Figure 3: The bold edges form an infinite circuit in G , the broken edges indicate the corresponding cut in G^*

Theorem 3.4. *Let G be a countable⁵ graph satisfying (*).*

- (i) G has a dual if and only if G is planar.
- (ii) If G^* is a dual of G , then G^* satisfies (*), G is a dual of G^* , and this is witnessed by the inverse bijection of *.
- (iii) If G^* is a dual of G and $F \subseteq E(G)$, then $F \in \mathcal{C}(\tilde{G})$ if and only if F^* is a cut in G^* .

We shall prove Theorem 3.4 in the next section.

Since all finitary duals of a 3-connected graph are again 3-connected [13], Theorems 3.3 and 3.4 (ii) together imply at once that the dual of a 3-connected graph is unique:

Corollary 3.5. *A 3-connected graph has at most one dual, up to isomorphism.*

4 Proof of the duality theorem

Recall that, by Theorem 3.1, any graph G with a finitary dual G^* satisfies (*). Our first aim is to show that if G^* is a dual of G , then G^* too satisfies (*).

We need two lemmas. The first we quote from [2]:

Lemma 4.1. [2] *Let G be a 2-connected graph satisfying (*), and let U be a finite set of vertices in G . Then we can contract edges of G so that no two vertices from U are identified, the graph H obtained has only finitely many edges and vertices, and every cut in H is also a cut of G .*

⁵By Lemma 4.4 below, the countability assumption is redundant for 2-connected graphs, and therefore inessential. If we agree to call a graph ‘planar’ as soon as it has neither a K_5 nor a $K_{3,3}$ minor, then Theorem 3.4 becomes true also for uncountable graphs.

Lemma 4.2. *Let G be a 2-connected graph satisfying $(*)$, and let C be an infinite circuit in \tilde{G} . Let X be a finite set of edges meeting C in exactly one edge e . Then there is a finite circuit in G that meets X precisely in e .*

Proof. Apply Lemma 4.1 to G , taking as U the set of endvertices of the edges in X . Consider the finite graph H returned by the lemma. Applying Theorem 2.2 twice, we deduce from $C \in \mathcal{C}(\tilde{G})$ and the separation property of H that $C \cap E(H) \in \mathcal{C}(H)$. Let $C' \subseteq C \cap E(H)$ be a circuit containing e . As C meets X only in e , so does C' . Since no two vertices from U were identified when G was contracted to H , the branch sets of the contraction induce no edge from X in G . We can therefore expand C' to a finite circuit in G that still meets X only in e . \square

We need the following strong version of Menger's theorem for countable graphs.

Theorem 4.3 (Aharoni [1]). *For any countable graph G and two sets A, B of vertices in G there exist a set \mathcal{P} of disjoint A - B paths and an A - B separator X in G such that X consists of a choice of one vertex from each of the paths in \mathcal{P} .*

As every uncountable connected graph has a vertex of uncountable degree, it is easy to show that an uncountable 2-connected graph contains two vertices joined by uncountably many independent paths. (See eg. [8, Lemma 2.1], or Thomassen [13].) Thus:

Lemma 4.4. *A 2-connected graph satisfying $(*)$ is countable.*

We can now show that, unlike with finitary duals, the class of graphs satisfying $(*)$ is closed under taking duals.

Lemma 4.5. *Any dual G^* of a graph G satisfies $(*)$.*

Proof. Suppose G^* violates $(*)$. Then there are two vertices in G^* , x and y say, that cannot be separated by finitely many edges. We may assume that x and y lie in the same block B^* of G^* . Let B be the subgraph of G consisting of the edges e for which $e^* \in E(B^*)$ together with their incident vertices.

Let e be an edge of B , and consider an edge f that lies in the same block of G as e . It is not hard to see that there is a finite circuit containing both e and f . Thus, by duality, e^* and f^* lie in a common bond F^* of G^* . The edges in F^* that lie in B^* suffice to separate the endvertices of e^* in G^* . By the minimality of F^* , these are all its edges, including f^* , and thus $f \in E(B)$. This shows that B is the union of blocks of G .

Since an edge set is a circuit (resp. bond) of a graph if and only if it is a circuit (resp. bond) in one of its blocks, we see that B^* is a dual of the graph B . We may therefore assume that G^* is 2-connected, and thus, by Lemma 4.4, countable.

By Theorem 4.3 applied to the line graph of G^* (which is countable because G^* is), we can find in G^* an infinite set \mathcal{P} of edge-disjoint x - y paths and a set C^* of edges separating x from y , such that C^* consists of a choice of one edge from each path in \mathcal{P} . Then C^* is a bond in G^* , and $C = \{e \mid e^* \in C^*\}$ is an infinite circuit in \tilde{G} .

Pick an edge $e \in C$. We claim the following:

There is an infinite sequence of distinct finite circuits C_1, C_2, \dots in G , each containing e , and such that $C_i \setminus C$ and $C_j \setminus C$ are non-empty and disjoint for all $i \neq j$. (1)

To prove (1), assume inductively that C_1, \dots, C_{i-1} have been constructed, and put

$$X := \{e\} \cup \bigcup_{j < i} C_j \setminus C.$$

Since C meets X only in e , Lemma 4.2 gives us a finite circuit C_i that contains e and does not meet $C_1 \cup \dots \cup C_{i-1}$ outside C . As both C and C_i are circuits in \tilde{G} , neither contains the other properly, so $C_i \setminus C \neq \emptyset$. This proves (1).

Let u, v be the endvertices of e^* in G^* . Each of the sets C_i^* is a cut in G^* that contains e^* , and hence separates u from v in G^* . Denote by P the path in \mathcal{P} that contains e^* . Since $E(P)$ meets C^* only in e^* , no edge of P other than e^* lies in more than one of the sets C_i^* (by (1)). Therefore only finitely many of the sets C_i^* meet $E(P - e^*)$; let C_n^* be one that does not. Since C_n^* is finite, there is a path $Q \in \mathcal{P}$ that has no edge in C_n^* . But then $(P - e^*) \cup Q$ is a connected subgraph of G^* that avoids C_n^* but contains both u and v , a contradiction. \square

As pointed out before, every dual of a graph is also its finitary dual, and we have just seen that it satisfies (*). Our next aim is to show that, conversely, every finitary dual satisfying (*) is even a dual. We need the following lemma:

Lemma 4.6. *A set $F \subseteq E(G)$ in a graph G is a cut if and only if it meets every finite circuit in an even number of edges.*

Proof. Clearly, a cut meets every finite circuit in an even number of edges, so let us prove the other direction.

Let G' be the graph obtained from G by contracting every edge not in F . Then G' is bipartite (in particular, loopless), since any odd circuit would give rise to a finite circuit in G meeting F in an odd number of edges. The bipartition of G' induces a partition (A, B) of the vertex set of G such that every edge in F has one vertex in A and the other in B and such that no edge outside F has that property. Thus, $F = E(A, B)$ is a cut. \square

Lemma 4.7. *Let G be a 2-connected graph, and let G^* be a finitary dual of G that satisfies (*). Then the following assertions hold:*

- (i) G^* is a dual of G (witnessed by the same map *).
- (ii) A set $F \subseteq E(G)$ lies in $\mathcal{C}(\tilde{G})$ if and only if F^* is a cut in G^* .

Proof. We first prove (ii). Let $F \in \mathcal{C}(\tilde{G})$ be given. To show that F^* is a cut in G^* , it suffices by Lemma 4.6 to show that F^* meets every finite circuit Z^* in G^* in an even number of edges. Since G is a finitary dual of G^* (Theorem 3.2), Z is a finite cut in G . Hence by Theorem 2.2, $|F \cap Z| = |F^* \cap Z^*|$ is even.

Similarly, consider a cut F^* in G^* . To show that $F \in \mathcal{C}(\tilde{G})$, it suffices by Theorem 2.2 to show that F meets every finite cut D of G evenly. But $D^* \in \mathcal{C}(\tilde{G}^*)$ since G is a finitary dual of G^* (Theorem 3.2). So $|F \cap D| = |F^* \cap D^*|$ is even by the trivial direction of Lemma 4.6 and the fact that D^* is a disjoint union of circuits.

To prove (i), let now C be a circuit in \tilde{G} ; we have to show that C^* is a bond in G^* . We have already shown that C^* is a cut in G^* , and that any cut $F^* \subseteq C^*$ corresponds to a set $F \in \mathcal{C}(G)$. Since F cannot be a proper subset of C unless it is empty, we deduce that C^* is a bond.

Conversely, let F^* be a bond in G^* . Then F is a minimal non-empty element of $\mathcal{C}(G)$. By Theorem 2.1, F must be a circuit. \square

For a proof of Theorem 3.4, it remains to combine Lemmas 4.5 and 4.7 with Thomassen's results on finitary duals, and to extend the result from 2-connected to arbitrary graphs.

The latter is standard for finite graphs, but we have to be more careful here. (Indeed, Lemma 4.8 below fails for finitary duals.) If G has a finitary dual G^* and B is a block of G , let B^* denote the subgraph of G^* formed by the edges e^* with $e \in B$ and their incident vertices.

Lemma 4.8. *Let a graph G have a finitary dual G^* that satisfies (*). If B is a block of G , then B^* is a block of G^* and a finitary dual of B . If G^* is even a dual of G then B^* is a dual of B .*

Proof. Two edges $e, f \in G$ lie in a common block of G if and only if they lie in a common finite circuit of \tilde{G} . For a proof that B^* is a block of G^* , it therefore suffices to show that the edges e^* and f^* lie in a common block of G^* if and only if they lie in a common finite bond of G^* .

If e^* and f^* lie in a common bond F^* of G^* , we proceed in a similar way as in the proof of Lemma 4.5.

Now suppose that e^* and f^* lie in a common block B^* of G^* . Then B^* has a finite circuit containing both e^* and f^* . Deleting e^* and f^* from this circuit, we obtain the edge sets of two paths, P and Q . Suppose that every $X \subseteq E(G^*)$ separating P and Q is infinite. Then there are also two vertices, one in P and the other in Q , that cannot be separated in G^* by finitely many edges. Consequently, we find infinitely many edge-disjoint paths connecting these two vertices, a contradiction to (*). Therefore, there is a finite set $F^* \subseteq E(G^*)$ separating P from Q in G^* , which clearly contains e^* and f^* . If we choose F^* to be minimal then it is a bond.

It remains to show that B^* is a finitary dual (resp. dual) of B . But since B^* is a block of G^* , a set $F^* \subseteq E(B^*)$ is a bond of B^* if and only if it is a bond of G^* . Similarly, a set $F \subseteq B$ is a circuit in \tilde{B} if and only if it is a circuit in \tilde{G} . The assertion therefore follows from the assumption that G^* is a finitary dual (resp. dual) of G . \square

Lemma 4.9. *If G and G^* are two graphs and $*$: $E(G) \rightarrow E(G^*)$ maps the blocks B of G to the blocks of G^* so that B^* is a dual of B , then G^* is a dual of G .*

Proof. It is easily checked that a subset of $E(G)$ is a circuit in \tilde{G} if and only if it is a circuit in \tilde{B} for some block B of G . Similarly, a subset of $E(G^*)$ is a bond in G^* if and only if it is a bond in some block of G^* . \square

Proof of Theorem 3.4. (i) By Lemmas 4.5, 4.8 and 4.9, G has a dual if and only if its blocks do. (To obtain a dual of G from duals of its blocks, take their disjoint union.) Similarly, a countable graph is planar if and only if its blocks are [10]. We may therefore assume that G is 2-connected.

If G is planar then, by Theorems 3.1 and 3.2, G has a finitary dual G^* that satisfies (*). By Lemma 4.7, G^* is even a dual of G . Conversely, if G has a dual G^* , then G is planar by Theorem 3.1.

(ii) Suppose that G^* is a dual of G . By Lemma 4.5, G^* satisfies (*). By Lemma 4.8, the subgraphs B^* of G^* , where B ranges over the blocks of G , are the blocks of G^* , and each B^* is a dual of B . We show that, conversely, B is a dual of B^* . Then, by Lemma 4.9, G is a dual of G^* .

By Theorem 3.2, B is a finitary dual of B^* . Now B satisfies (*), because G does so by assumption. Hence by Lemma 4.7, B is a dual of B^* .

(iii) For 2-connected graphs, this is Lemma 4.7 (ii). The general case reduces easily to this with the help of Lemma 4.6 and Theorem 2.2. \square

5 Locally finite duals

We started out by observing that a dual of a locally finite graph may have vertices of infinite degree. This raises the question under what circumstances the dual is locally finite. For 3-connected graphs, Thomassen gave the following characterisation in terms of *peripheral* circuits, circuits C whose incident vertices do not separate the graph and do not span any edges not in C .

Theorem 5.1 (Thomassen [12]). *Let G be a locally finite 3-connected graph. Then G has a locally finite finitary dual if and only if G is planar and every edge lies in exactly two finite peripheral circuits.*

Since locally finite graphs trivially satisfy condition (*), Lemma 4.7 implies that Theorem 5.1 still holds if the word ‘finitary’ is dropped.

To obtain another characterisation, we need the following extension of Tutte’s planarity criterion to locally finite graphs:

Theorem 5.2. [3] *Let G be a locally finite 3-connected graph. If G is planar then every edge appears in exactly two peripheral circuits. Conversely, if every edge appears in at most two peripheral circuits then G is planar.*

Theorem 5.3. *A locally finite 3-connected graph has a locally finite dual if and only if it is planar and all its peripheral circuits are finite.*

Proof. Let G be a locally finite 3-connected graph. If G has a locally finite dual then, by Theorem 5.1, G is planar and every edge lies in exactly two finite peripheral circuits. By Theorem 5.2, its edges cannot lie in any other peripheral circuits, so all peripheral circuits are finite.

Conversely, if G is planar and all its peripheral circuits are finite then, by Theorems 5.2 and 5.1, G has a locally finite finitary dual. By Lemma 4.7, this is in fact a dual. \square

6 Duality in terms of spanning trees

In this section we show that our notion of duality permits the extension of another well-known duality theorem for finite graphs: that the complement of the edge set of any spanning tree of G defines a spanning tree in any dual of G , and conversely that any two graphs whose edge sets are in bijective correspondence

so that their spanning trees complement each other as above form a pair of duals.

It is not difficult to see that the verbatim analogue of this fails for infinite graphs. Indeed, the edge set of an ordinary spanning tree of G might contain an infinite circuit C (such as the edges of the double ray D in Figure 1), in which case C^* would be a cut in G^* , and $G^* - C^*$ could not contain a spanning tree of G^* . However, the following adjustment to the notion of a spanning tree makes an extension possible.

Let us call a spanning tree T of G *acirclic* (under ITOP) if its closure in \tilde{G} contains no circle – or equivalently, if its edges contain no circuit of \tilde{G} . (We remark that if G is locally finite then its acirclic spanning trees are precisely its end-faithful spanning trees [9].)⁶

Theorem 6.1. *Let $G = (V, E)$ and $G^* = (V^*, E^*)$ be connected⁷ graphs satisfying (*), and let $*$: $E \rightarrow E^*$ be a bijection. Then the following two assertions are equivalent:*

- (i) *G and G^* are duals of each other, and this is witnessed by the map $*$ and its inverse.*
- (ii) *Given a set $F \subseteq E$, the graph (V, F) is an acirclic spanning tree of G if and only if $(V^*, E^* \setminus F^*)$ is an acirclic spanning tree of G^* (both in ITOP).*

Before we prove Theorem 6.1, let us show that those acirclic spanning trees always exist. We need the following easy fact, whose proof is the same as for finite graphs [7, Lemma 1.9.4].

Lemma 6.2. *Every cut in a graph is a disjoint union of bonds.*

Theorem 6.3 below settles Problem 7.9 of [9].

Theorem 6.3. *Every connected graph G satisfying (*) has a spanning tree whose closure in \tilde{G} contains no circle.*

Proof. We may assume that $G = (V, E)$ is 2-connected, since the union T of acirclic spanning trees of the blocks of G is always an acirclic spanning tree of G . (Indeed, any circle C in the closure of T must contain two edges e, e' from different blocks; if x is a cutvertex separating these blocks in G , then \hat{e} and \hat{e}' are separated topologically in the space $\tilde{G} - x$, which contradicts the connectedness of the open arc $C - x$.)

By Lemma 4.4, E has an enumeration e_1, e_2, \dots . Put $S_0 = T_0 = E$, and inductively for $n = 1, 2, \dots$ define $S_n, T_n \subseteq E$ as follows. Given n , denote by i_n the least index i such that \tilde{G} has a circuit $C \subseteq S_{n-1}$ that contains both e_n and e_i ; if there is no such circuit, let $i_n = \infty$. Analogously, choose j_n minimum so that some bond $B \subseteq T_{n-1}$ contains both e_n and e_{j_n} ; if there is no such bond, let $j_n = \infty$. If $i_n < j_n$ put $S_n := S_{n-1} - e_n$ and $T_n := T_{n-1}$; if $i_n > j_n$ put

⁶In [9], we considered the more general concept of ‘topological spanning trees’. These are acirclic and path-connected subspaces of $|G|$ or \tilde{G} , which however need not induce a connected subgraph of G . Our acirclic spanning trees defined above *are* meant to be connected as graphs: they are just graph-theoretical spanning trees of G with the additional property that their closure in \tilde{G} contains no circle (and is therefore a topological spanning tree of \tilde{G}).

⁷This assumption is for convenience only. For disconnected graphs, one has to replace ‘acirclic spanning tree’ with ‘subgraph inducing an acirclic spanning tree in every component’.

$S_n := S_{n-1}$ and $T_n := T_{n-1} - e_n$; if $i_n = j_n$, choose arbitrarily whether to delete e_n from S_{n-1} or from T_{n-1} . (The ambiguity here is deliberate, to keep the definition symmetrical in S and T . This symmetry will be used later.) Then $S := \bigcap_{n=1}^{\infty} S_n$ and $T := \bigcap_{n=1}^{\infty} T_n$ partition E . More precisely:

$$\text{For all } n \leq m, \text{ the edge } e_n \text{ lies in exactly one of the two sets } S_m, T_m. \quad (2)$$

We shall prove that S contains no circuit, and that T contains no bond. Then (V, S) is an acirclic spanning tree, completing the proof. (Indeed, if (V, S) is not connected there is a bond B in G such that $B \subseteq E \setminus S = T$, a contradiction.)

So, assume that S contains a circuit or that T contains a bond, and choose i minimum so that there is a set $C \subseteq E$ with $e_i \in C$, and such that $C \subseteq S$ is a circuit in \tilde{G} or such that $C \subseteq T$ is a bond of G . We first assume that $C \subseteq S$, i.e. that C is a circuit.

If C is finite, consider the edge $e_k \in C$ with k maximum. As $e_k \in C \subseteq S \subseteq S_k$, we have $e_k \notin T_k$ by (2). Then $j_k \leq i < \infty$, so there is a bond $D \subseteq T_{k-1}$ containing e_k . As $e_k \in C \cap D$, Lemma 4.6 implies that $C \cap D$ contains another edge e_j , with $j \leq k-1$ by the choice of k . As $e_j \in C \cap D \subseteq S_{k-1} \cap T_{k-1}$, this contradicts (2).

Therefore C is infinite. Let $e_i, e_{k_1}, e_{k_2}, \dots$ be distinct edges in C , and note that $i < k_l$ for each l , by the choice of i . Since S_{k_l-1} contains the circuit $C \ni e_{k_l}$ but $e_{k_l} \in S_{k_l}$, there is a bond $D_l \subseteq T_{k_l-1}$ containing e_{k_l} and an edge e_{m_l} with $m_l \leq i$; otherwise we would have deleted e_{k_l} from S_{k_l-1} to obtain S_{k_l} . Since $e_i \in S_{k_l-1}$ implies $e_i \notin T_{k_l-1}$, by (2), we cannot have $m_l = i$, so in fact $m_l < i$. Choose $m < i$ so that $m = m_l$ for infinitely many l , and let L_0 denote the set of these l . Thus, $e_m \in D_l$ for every $l \in L_0$.

Put $E^n := \{e_1, \dots, e_n\}$. Inductively for $n = 1, 2, \dots$, choose infinite index sets $L_0 \supseteq L_1 \supseteq \dots$ so that, for each n , the sets $D^n := D_l \cap E^n$ coincide for all $l \in L_n$. We claim that $D = \bigcup_{n=1}^{\infty} D^n$ is a cut of G contained in T . The edge $e_m \in D$ will then lie in some bond $B \subseteq T$ (Lemma 6.2), which contradicts the minimal choice of i as $m < i$.

To show that D is a cut, it suffices to check that D meets every finite circuit C' in an even number of edges (Lemma 4.6). Choose n large enough that $C' \subseteq E^n$. Then

$$C' \cap D = C' \cap E^n \cap D = C' \cap D_l$$

for every $l \in L_n$. But $|C' \cap D_l|$ is even since D_l is a cut, again by Lemma 4.6.

To show that $D \subseteq T$, consider any edge $e_n \in D$. Then $e_n \in D_l$ for every $l \in L_n$. By definition, D_l is a subset of T_{k_l-1} . Now as L_n is infinite, we may assume that $k_l > n$. But then $e_n \in T_{k_l-1}$ implies that $e_n \in T$, as desired.

Finally, if $C \subseteq T$ is a bond rather than a circuit, the proof is analogous to the above, with the roles of S and T and of circuits and bonds interchanged. Instead of Lemmas 4.6 and 6.2 we use Theorems 2.2 and 2.1. \square

Proof of Theorem 6.1. (i) \Rightarrow (ii): Let $T = (V, F)$ be an acirclic spanning tree of G in ITOP. Then $E^* \setminus F^*$ contains no circuit C^* of \tilde{G}^* , since C would then be a cut of G missed by T . Similarly $(V^*, E^* \setminus F^*)$ must be connected: if not, then F^* contains a bond of G^* , and F contains the corresponding circuit of \tilde{G} . The converse implication follows by symmetry, since G is a dual of G^* .

(ii) \Rightarrow (i): We show that the map $*$ makes G^* a finitary dual of G . Then Lemma 4.8 implies that for each block B of G the block B^* of G^* is a finitary

dual. Then B^* is even a dual of B (with the same map $*$, by Lemma 4.7), and by Theorem 3.4 the inverse of $*$ makes B a dual of B^* . With Lemma 4.9 we obtain (i).

So consider a finite circuit C of G . We first show that C^* contains a cut of G^* . Using Theorem 6.3, choose an acirclic spanning tree S^* of G^* . If C^* contains no cut, we can join up the components of $S^* - C^*$ by finitely many edges from $E^* \setminus C^*$ to form another spanning tree T^* of G^* . Then T^* , too, is acirclic: any circle in its closure contains an arc A^* that contains infinitely many edges but avoids the (finitely many) new edges and hence lies in the closure of S^* , so the union of A^* with a suitable path from S^* contains a circle in the closure of S^* . Now use (ii) to find an acirclic spanning tree T of G corresponding to T^* . Since T^* contains no edge from C^* , the edges of T include C , a contradiction.

To show that C^* is even a minimal cut in G^* , we show that for every $e \in C$ and $A := C \setminus \{e\}$ the graph $G^* - A^*$ is connected. To do so, it suffices to find a spanning tree T^* of G^* with no edge in A^* , and hence by (ii) to find an acirclic spanning tree of G whose edges include A . Let S be any acirclic spanning tree of G . Since A is finite but contains no circuit, we can obtain another spanning tree T from S by adding all the edges from A and deleting some (finitely many) edges not in A . As before, T is acirclic in \tilde{G} because S was, and hence is as desired.

It remains to show that if B^* is a finite bond in G^* then B is a circuit in G . As before, we first show that B contains a circuit. If not, we can modify an acirclic spanning tree of G into one whose edges include B , which by (ii) corresponds to a spanning tree of G^* that has no edge in B^* (contradiction). On the other hand, given any proper subset D^* of B^* , we can modify an acirclic spanning tree of G^* into one missing D^* , because D^* contains no cut of G^* . Then this tree corresponds by (ii) to a spanning tree of G whose edges include D , so D is not a circuit in G . \square

7 Colouring-flow duality and circuit covers

As an application of Theorem 3.4 and our results from Section 4, we now show that the edge set of every bridgeless locally finite planar graph can be covered by two elements of its cycle space. For finite graphs, this is a well-known reformulation of the four colour theorem. For infinite graphs, of course, it must fail as long as the cycle space contains only finite sets of edges.

In our setting, however, Theorem 3.4 enables us to imitate the finite result (and its proof from the four colour theorem), because 4-colourability extends by compactness [4]. Rather than assuming that G is locally finite, we work slightly more generally in \tilde{G} .

Theorem 7.1. *Let G be a bridgeless planar graph satisfying $(*)$. Then there are $Z_1, Z_2 \in \mathcal{C}(\tilde{G})$ such that $E(G) = Z_1 \cup Z_2$.*

Proof. Assume that we find for every block B of G elements Z_1^B, Z_2^B of the cycle space of B such that $E(B) = Z_1^B \cup Z_2^B$. From Theorem 2.2 follows that for $i = 1, 2$, $Z_i^B \in \mathcal{C}(G)$ and then also $Z_i := \sum_B Z_i^B \in \mathcal{C}(G)$, where the sum ranges over the blocks of G . Clearly, we get $E(G) = Z_1 \cup Z_2$. As G is bridgeless, we may therefore assume that G is 2-connected.

By Theorem 3.4, G has a dual G^* and is itself a dual of G^* , which therefore is planar too. By the four colour theorem and compactness [4], G^* has chromatic number at most 4. Choose a 4-colouring $c: V(G^*) \rightarrow \mathbb{Z}_2 \times \mathbb{Z}_2$ of G^* . For $i = 1, 2$, let $c_i: V(G^*) \rightarrow \mathbb{Z}_2$ be c followed by the projection to the i th coordinate, define $f_i: E(G) \rightarrow \mathbb{Z}_2$ by $f_i(e) := c_i(v) + c_i(w)$ where v and w are the endvertices of e^* , and put $Z_i := f_i^{-1}(1)$.

Let us show that every edge e of G lies in Z_1 or Z_2 . If not, then $f_1(e) = f_2(e) = 0$, and hence $c(v) = c(w)$ for $e^* =: vw$. But this contradicts our assumption that c is a proper colouring of G^* .

Next we show that $Z_i \in \mathcal{C}(\tilde{G})$, for both $i = 1, 2$. By Theorem 2.2, it suffices to show that Z_i meets every finite cut F of G in an even number of edges, ie. that

$$f_i(F) := \sum_{e \in F} f_i(e) = 0.$$

As every cut is a disjoint union of bonds (Lemma 6.2), we may assume that F is a bond. Then F^* is a circuit in G^* . Hence,

$$f_i(F) = \sum_{e^* = vw \in F^*} (c_i(v) + c_i(w)) = 2 \sum_{u \in U} c_i(u) = 0,$$

where U is the vertex set of the cycle in G^* whose edge set is F^* . □

References

- [1] R. Aharoni, *Menger's theorem for countable graphs*, J. Combin. Theory (Series B) **43** (1987), 303–313.
- [2] H. Bruhn, R. Diestel, and M. Stein, *Cycle-cocycle partitions and faithful cycle covers for locally finite graphs*, J. Graph Theory **50** (2005), 150–161.
- [3] H. Bruhn and M. Stein, *MacLane's planarity criterion for locally finite graphs*, To appear in *J. Combin. Theory (Series B)*.
- [4] N.G. de Bruijn and P. Erdős, *A colour problem for infinite graphs and a problem in the theory of relations*, Indag. Math. **13** (1951), 371–373.
- [5] R. Diestel, *End spaces and spanning trees*, To appear in *J. Combin. Theory (Series B)*.
- [6] ———, *The cycle space of an infinite graph*, Comb., Probab. Comput. **14** (2005), 59–79.
- [7] ———, *Graph theory* (3rd edition), Springer-Verlag, 2005.
- [8] R. Diestel and D. Kühn, *Graph-theoretical versus topological ends of graphs*, J. Combin. Theory (Series B) **87** (2003), 197–206.
- [9] ———, *Topological paths, cycles and spanning trees in infinite graphs*, Europ. J. Combinatorics **25** (2004), 835–862.
- [10] G.A. Dirac and S. Schuster, *A theorem of Kuratowski*, Nederl. Akad. Wetensch. Proc. Ser. A **57** (1954), 343–348.

- [11] H. Freudenthal, *Neuaufbau der Endentheorie*, Annals of Mathematics **43** (1942), 261–279.
- [12] C. Thomassen, *Planarity and duality of finite and infinite graphs*, J. Combin. Theory (Series B) **29** (1980), 244–271.
- [13] ———, *Duality of infinite graphs*, J. Combin. Theory (Series B) **33** (1982), 137–160.
- [14] H. Whitney, *Non-separable and planar graphs*, Trans. Am. Math. Soc. **34** (1932), 339–362.

Version 18 Okt 2005

Henning Bruhn <hbruhn@gmx.net>
Reinhard Diestel <diestel@math.uni-hamburg.de>
Mathematisches Seminar
Universität Hamburg
Bundesstraße 55
20146 Hamburg
Germany