

TOPOLOGICAL CYCLE MATROIDS OF INFINITE GRAPHS

Johannes Carmesin

September 28, 2016

Abstract

We prove that the topological cycles of an arbitrary infinite graph together with its topological ends form a matroid. This matroid is, in general, neither finitary nor cofinitary.

1 Introduction

Many theorems about finite graphs and their cycles do not extend to infinite graphs and their finite cycles. However, many such theorems do extend to locally finite graphs together with their topological cycles, see for example [7, 17, 18, 21], and [13] for a survey. These *topological cycles* are homeomorphic images of the unit circle in the topological space obtained from the graph by adding certain points at infinity called ends.

Bruhn and Diestel gave an explanation why many of these theorems extended: the topological cycles of a locally finite graph form a matroid [8]. This matroidal point of view allowed for new proof techniques and abstracting the topological properties of the topological cycles often led to simpler proofs. For non-locally finite graphs various notions of end boundaries have been suggested [13], each of which gives rise to its own notion of topological cycles.

To compare these end boundaries we will not refer directly to topology but instead compare the matroids they induce. However for some of these notions, the matroids have finite circuits which are not finite cycles of the graph. A consequence of this is that there are non-isomorphic (3-connected) graphs inducing isomorphic matroids. For others we even do not always get a matroid.

Here we show that the topological end boundary, which had not been considered for this purpose before, lacks these defects. More precisely:

Theorem 1.1. *For any graph G , the topological cycles of G together with its topological ends form a matroid.*

Moreover, for non-isomorphic 3-connected graphs, these matroids are non-isomorphic.

Furthermore, all matroids that arise as cycle matroids for one of the other boundaries are minors of these cycle matroids. For the one boundary, where the topological cycles do not induce a matroid for all graphs, Theorem 1.1 implies a characterisation when they do. The various notions of boundary and the corresponding characterisations are compared in Subsection 1.1.

The question whether the topological cycles in the graph G together with the boundary B induce a matroid is closely related to the question whether G has a spanning tree whose ends are equal to B . Indeed, any such spanning tree is an example of a base in the topological cycle matroid.

In our proof we use a result of [10] which ensures the existence of such spanning trees for the topological ends. We then combine this with the theory of trees of matroids [4].

1.1 Comparing the end boundaries

Bruhn and Diestel showed that the dual of the finite-bond matroid of a graph G is given by the topological cycles of G together with its edge ends [8]. However, after deleting parallel edges, any component of such a matroid is countable.

Hence in order to construct matroids that are nontrivially uncountable, we have to consider topological cycles of different topological spaces. One such space is VTOP, which is obtained from the graph by adding the vertex ends. In Figure 1, we depicted a graph whose topological cycles in VTOP do not induce a matroid.

The reason why this example works is that the topological cycle C goes through a dominated vertex end. Here a vertex v *dominates* a vertex end ω if there is an infinite v -fan to some ray belonging to ω .

One way to ‘repair’ VTOP is to identify each vertex ends with the vertices dominating it. The resulting space is called ITOP. A consequence of Theorem 1.1 is the following.

Corollary 1.2. *For any graph, the topological cycles of ITOP form a matroid.*

The matroids we get from Corollary 1.2 are more complicated than the ones for the edge ends in the sense that they are not always cofinitary.




Figure 1: The dominated ladder is obtained from the one ended ladder by adding a vertex that is adjacent to every vertex on the upper side of the ladder. The topological cycles of VTOP of the dominated ladder do not induce a matroid as they violate the elimination axiom (C3): We cannot eliminate all the triangles from the grey cycle C .

However, there are still non-isomorphic 3-connected graphs whose ITOP-matroids are isomorphic.

Another way to ‘repair’ VTOP is to delete the dominated vertex ends. Diestel and Kühn [16] showed that the remaining vertex ends are given by the topological ends, and in this case the topological cycles induce a matroid by Theorem 1.1.

In 1969, Higgs proved that the set of finite cycles and double rays of a graph G is the set of circuits of a matroid if and only if G does not have a subdivision of the Bean-graph [20]. Using Theorem 1.1, we get a result for the topological cycles of VTOP in the same spirit.

Corollary 1.3. *The topological cycles of VTOP induce a matroid if and only if G does not have a subdivision of the dominated ladder, which is depicted in Figure 1.*

Theorem 1.1 extends to ‘Psi-Matroids’: Given a set Ψ of topological ends, let C_Ψ be the set of those topological cycles that only use topological ends in Ψ . Let D_Ψ be the set of those bonds that have no topological end of Ψ in their closure. We prove the following strengthening of Theorem 1.1.

Theorem 1.4. *Let Ψ be a Borel set of topological ends. Then C_Ψ and D_Ψ are the sets of circuits and cocircuits of a matroid.*

If we leave out the assumption that Ψ is Borel, then this theorem becomes false, see [3] for details. We also can extend the main result of [3], see Section 5 for details.

The proof of Theorem 1.4 uses two different tools: the tree-decompositions constructed in [10], and the theory of trees of matroids from [4]. These tools and some basic notions are explained in Section 2. After some intermediate results in Section 3, we prove Theorem 1.4 in Section 4. Then, in Section 5 we deduce from it the other theorems mentioned in the Introduction. In the Discussion at the end, we mention an open problem about the class of graphic matroids.

2 Preliminaries

Throughout, notation and terminology for graphs are that of [15] unless defined differently. G always denotes a graph. We denote the complement of a set X by X^c . Throughout this paper, *even* always means finite and a multiple of 2. An edge set F in a graph is a *cut* if there is a partition of the set of vertices such that F is the set of edges with precisely one endvertex in each partition class. A vertex set *covers* a cut if every edge of the cut is incident with a vertex of that set. A cut is *finitely coverable* if there is a finite vertex set covering it. A *bond* is a minimal nonempty cut.

For us, a *separation* is just an edge set. The *boundary* $\partial(X)$ of a separation X is the set of those vertices adjacent with an edge from X and one from X^c . The *order* of X is the cardinality of $\partial(X)$. Given a connected subgraph C of G , its *incidence set* s_C consists of those edges incident with at least one vertex of C .

Given a separation X of finite order and a vertex end ω , then there is a unique component C of $G - \partial(X)$ in which ω lives. We say that ω *lives* in X if $s_C \subseteq X$.

2.1 Tree-decompositions and infinite graphs

A *tree-decomposition* of G consists of a tree T together with a family of subgraphs $(P_t \mid t \in V(T))$ of G such that every vertex and edge of G is in at least one of these subgraphs, and such that if v is a vertex of both P_t and P_w , then it is a vertex of each P_u , where u lies on the t - w -path in T . Moreover, each edge of G is contained in precisely one P_t . The subgraphs P_t are the *parts* of the tree-decomposition. Two parts P_v and P_u are *adjacent* if v and u are adjacent in the decomposition tree T . Sometimes, the ‘Moreover’-part is not part of the definition of tree-decomposition. However, both these two definitions give the same concept of tree-decomposition since any tree-decomposition without this additional property can easily be changed to one with this property by deleting edges from the parts appropriately.

The *adhesion of a tree-decomposition is finite* if any two (adjacent) parts intersect finitely.

In this paper we will always work with a *rooted* tree-decomposition, which simply means that the decomposition tree is rooted. The choice of a root, endows the node set of the tree with a partial ordering \leq with the root being the smallest node: $s < t$ if t and the root are in different components of $T - s$. In this paper we follow the convention that in a rooted tree all its edges are directed away from the root, so from the \leq -smaller endvertex to the \leq -bigger one. If st is a directed edge, then it is directed from s to t . Given a directed edge tu of a rooted decomposition tree T of a tree-decomposition $(T, (P_t \mid t \in V(T)))$, the *separation $S[tu]$ corresponding to tu* is the set of those edges which are in parts P_w with $w \geq u$. The set $V[u]$ of *vertices above u* is $\bigcup_{w \geq u} V(P_w)$.

A rooted tree-decomposition $(T, (P_t \mid t \in V(T)))$ is *strongly exhausting* if it satisfies the following:

1. if st and tu are directed edges, then the separators $\partial(S[st])$ and $\partial(S[tu])$ corresponding to st and tu , respectively, are vertex-disjoint; AND
2. for each directed edge tu of T the graph $(V[u], S[tu])$ is connected.

Note that in any strongly exhausting tree-decomposition, the following strengthening of 1 also holds: if st and uw are directed edges with $t \leq u$, then the separators $\partial(S[st])$ and $\partial(S[uw])$ corresponding to st and uw , respectively, are vertex-disjoint. We remark that any tree-decomposition can easily be modified to one satisfying 2.

Given a tree-decomposition $(T, (P_t \mid t \in V(T)))$ of finite adhesion of a graph G , a vertex end ω of G *lives* in an end μ of T if for every directed edge tu of T such that μ lives in the component of $T - tu$ containing u , the vertex end ω lives in $S[tu]$. The ends of T *define precisely* the topological ends of G if

- in every end of T lives a unique vertex end of G and it is topological; AND
- every topological end of G lives in an end of T .

A key tool in our proof is the following main result of [10]:

Theorem 2.1 ([10, Theorem 1 and Remark 6.6]). *Every graph G has a rooted tree-decomposition $(T, (P_t \mid t \in V(T)))$ of finite adhesion such that the ends of T define precisely the topological ends of G .*

Moreover, $(T, (P_t \mid t \in V(T)))$ is strongly exhausting.

Given a part P_t of a tree-decomposition, the *torso* H_t is the multigraph with vertex set $V(P_t)$ whose edge set is the disjoint union of the edge set of P_t together with the edge set of the complete graph with vertex set $V(P_t) \cap V(P_u)$ for each neighbour u of t in the tree.

We denote the set of vertex ends of a graph G by $\Omega(G)$. A vertex v is in the *closure* of an edge set F if there is an infinite fan from v to $V(F)$. A vertex end ω is in the *closure* of an edge set F if every finite order separation X in which ω lives meets F . It is straightforward to show that a vertex end ω is in the closure of an edge set F if and only if every ray (equivalently: some ray) belonging to ω cannot be separated from F by removing finitely many vertices. A vertex end ω *lives in* a component C if it is in the closure of the incidence set s_C . A *comb* is a subdivision of the graph obtained from the ray by attaching a leaf at each of its vertices. These newly added vertices are the *teeth* of the comb. The Star-Comb-Lemma is the following.

Lemma 2.2 (Diestel [14, Lemma 1.2]). *Let U be an infinite set of vertices in a connected graph G . Then either there is a comb with all its teeth in U or a subdivision of the infinite star S with all leaves in U .*

Corollary 2.3. *Every infinite edge set has a vertex end or a vertex in its closure.* □

2.2 Infinite matroids

An excellent introduction to infinite matroids is in [9]. Here we rely on a characterisation of infinite matroids from [3], Theorem 2.4 below. Before we can state this characterisation, we list the axioms that appear in this characterisation. Let \mathcal{C} and \mathcal{D} be sets of subsets of a set E , which can be thought of as the sets of circuits and cocircuits of some matroid with groundset E , respectively.

(C1) The empty set is not in \mathcal{C} .

(C2) No element of \mathcal{C} is a subset of another.

(O1) $|C \cap D| \neq 1$ for all $C \in \mathcal{C}$ and $D \in \mathcal{D}$.

(O2) For all partitions $E = P \dot{\cup} Q \dot{\cup} \{e\}$ either $P + e$ includes an element of \mathcal{C} through e or $Q + e$ includes an element of \mathcal{D} through e .

We follow the convention that if we put a $*$ at an axiom A then this refers to the axiom obtained from A by replacing \mathcal{C} by \mathcal{D} , for example (C1 $*$) refers

to the axiom that the empty set is not in \mathcal{D} . A set $I \subseteq E$ is *independent* if it does not include any nonempty element of \mathcal{C} . Given $X \subseteq E$, a *base* of X is a maximal independent subset Y of X .

(IM) Given an independent set I and a superset X , there exists a base of X including I .

The proof of [3, Theorem 4.2] also proves the following:

Theorem 2.4. *Let E be some set and let $\mathcal{C}, \mathcal{D} \subseteq \mathcal{P}(E)$. Then there is a matroid M whose set of circuits is \mathcal{C} and whose set of cocircuits is \mathcal{D} if and only if \mathcal{C} and \mathcal{D} satisfy (C1), (C1*), (C2), (C2*), (O1), (O2), and (IM).*

Theorem 2.4 shows that the above axioms give an alternative axiomatisation of infinite matroids, which we use in this paper as a definition of infinite matroids. We call elements of \mathcal{C} *circuits* and elements of \mathcal{D} *cocircuits*. The *dual* of $(\mathcal{C}, \mathcal{D})$ is the matroid whose set of circuits is \mathcal{D} and whose set of cocircuits is \mathcal{C} .

A matroid $(\mathcal{C}, \mathcal{D})$ is *finitary* if every element of \mathcal{C} is finite, and it is *tame* if each element of \mathcal{C} intersects each element of \mathcal{D} only finitely. An example of a finitary matroid is the *finite-cycle matroid* of a graph G whose circuits are the edge sets of finite cycles of G and whose cocircuits are the bonds of G . We shall need the following lemma.

Lemma 2.5 ([6, Lemma 2.7]). *Suppose that M is a matroid, and \mathcal{C}, \mathcal{D} are collections of subsets of $E(M)$ such that \mathcal{C} contains every circuit of M , \mathcal{D} contains every cocircuit of M , and for every $o \in \mathcal{C}$, $b \in \mathcal{D}$, $|o \cap b| \neq 1$. Then the set of minimal nonempty elements of \mathcal{C} is the set of circuits of M and the set of minimal nonempty elements of \mathcal{D} is the set of cocircuits of M .*

2.3 Trees of presentations

In this subsection, we give simpler versions of the definitions of [4]. Binariness in general infinite matroids was studied by Christian [11]. However, the equivalent characterisations of finite binary matroids are not true in this general setting. Nevertheless, these characterisations extend to tame matroids [2]. In this paper all matroids are tame and we define: a tame matroid is *binary* if every circuit and cocircuit always intersect in an even number of edges. Roughly, a binary presentation of a tame matroid M is something like a pair of representations over \mathbb{F}_2 , one of M and of the dual of M , formally:

Definition 2.6. Let E be any set. A *binary presentation* Π on E consists of a pair (V, W) of sets of subsets of E satisfying (O2) and are orthogonal, that is, every $v \in V$ intersects any $w \in W$ evenly (so in particular finitely). We will sometimes denote the first element of Π by V_Π and the second by W_Π . We say that Π *presents* the matroid M if the circuits of M are the minimal nonempty elements of V_Π and the cocircuits of M are the minimal nonempty elements of W_Π .

Given a finitary binary matroid M , let \overline{V}_M be the set of those finite edge sets meeting each cocircuit evenly, and let \overline{W}_M be the set of those (finite or infinite) edge sets meeting each circuit evenly. Then $(\overline{V}_M, \overline{W}_M)$ is called the *canonical presentation* of M .

Definition 2.7. A *tree of binary presentations* \mathcal{T} consists of a tree T , together with functions \overline{V} and \overline{W} assigning to each node t of T a binary presentation $\Pi(t) = (\overline{V}(t), \overline{W}(t))$ on the ground set $E(t)$, such that for any two nodes t and t' of T , if $E(t) \cap E(t')$ is nonempty then tt' is an edge of T .

For any edge tt' of T we set $E(tt') = E(t) \cap E(t')$. We also define the *ground set* of \mathcal{T} to be $E = E(\mathcal{T}) = \left(\bigcup_{t \in V(T)} E(t) \right) \setminus \left(\bigcup_{tt' \in E(T)} E(tt') \right)$.

We shall refer to the edges which appear in some $E(t)$ but not in E as *virtual edges* of $M(t)$: thus the set of such virtual edges is $\bigcup_{tt' \in E(T)} E(tt')$.

A tree of binary presentations is a *tree of binary finitary presentations* if each presentation $\Pi(t)$ is a canonical presentation of some binary finitary matroid.

Remark 2.8. (*Motivation*) *The aim of this subsection is to state Theorem 2.10 below, which gives a criterion, when a binary tree of presentations $(T, \overline{V}, \overline{W})$ can be ‘glued together into a matroid’ with a binary presentation (V, W) . We have some freedom how to glue the matroids at infinity, which is given by a set Ψ of ends of T . The elements of V will be called ‘ Ψ -vectors’.*

Very roughly, they are given by picking an element of $\overline{V}(t)$ for each $t \in V(T)$ such that if ut is an edge, then the elements picked at $\overline{V}(t)$ and $\overline{V}(u)$ should coincide at the ‘overlap’ $E(tu)$. We also require that every end in the closure of the nodes where we picked nonempty elements is in Ψ . The Ψ -vectors consists of all non-virtual edges that appear in any of its picked elements. The elements of W are defined the same way but with \overline{W} in place of \overline{V} and Ψ^c in place of Ψ .

Definition 2.9. Let $\mathcal{T} = (T, \overline{V}, \overline{W})$ be a tree of binary presentations. A *pre-vector* of \mathcal{T} is a pair (S, \overline{v}) , where S is a subtree of T and \overline{v} is a function sending each node t of S to some $\overline{v}(t) \in \overline{V}(t)$, such that for each $t \in S$

we have $\bar{v}(t) \cap E(tu) = \bar{v}(u) \cap E(tu) \neq \emptyset$ if $u \in S$, and $\bar{v}(t) \cap E(tu) = \emptyset$ otherwise.

The *underlying vector* of (S, \bar{v}) is the set of those non-virtual edges in some $\bar{v}(t)$ for some $t \in V(S)$.

Now let Ψ be a set of vertex ends of T . A pre-vector (S, \bar{v}) is a Ψ -*pre-vector* if all vertex ends of S are in Ψ .

A Ψ -*vectors* is a symmetric difference of finitely many underlying vectors of Ψ -pre-vectors. We denote the space of all Ψ -vectors by $V_\Psi(\mathcal{T})$.

The following definitions of covectors and pre-covectors are almost identical except they use ' $\bar{W}(t)$ ' in place of ' $\bar{V}(t)$ '. *Pre-covectors* are defined like pre-vectors with ' $\bar{W}(t)$ ' in place of ' $\bar{V}(t)$ '. *Underlying covectors* are defined similar to underlying vectors. A pre-covector (S, \bar{w}) is a Ψ -*pre-covector* if all vertex ends of S are in Ψ .

The space $W_\Psi(\mathcal{T})$ of $\Psi^{\mathbb{G}}$ -*covectors* consists of those sets that are a symmetric differences of finitely many underlying covectors of $\Psi^{\mathbb{G}}$ -pre-covectors.

Finally, $\Pi_\Psi(\mathcal{T})$ is the pair $(V_\Psi(\mathcal{T}), W_\Psi(\mathcal{T}))$. The following is a consequence of the main result of [4, Theorem 8.3, and Lemma 6.8].

Theorem 2.10 ([4]). *Let $\mathcal{T} = (T, \bar{V}, \bar{W})$ be a tree of binary finitary presentations and Ψ a Borel set¹ of vertex ends of T , then $\Pi_\Psi(\mathcal{T})$ presents a binary matroid. Moreover, the sets of Ψ -vectors and $\Psi^{\mathbb{G}}$ -covectors satisfy (O1), (O2) and tameness.*

We shall also need the following related lemma, which is a combination of Lemma 6.6 and Lemma 6.8 from [4].

Lemma 2.11 ([4]). *Let $\mathcal{T} = (T, M)$ be a tree of binary finitary presentations and Ψ be any set of vertex ends of T . Any Ψ -vectors of \mathcal{T} and any $\Psi^{\mathbb{G}}$ -covectors of \mathcal{T} are orthogonal.*

3 Ends of graphs

Bruhn and Diestel [8] showed that the topological cycle matroid of a locally finite graph is the dual of the finite-bond matroid. So the whole topological information about the topological cycles is encoded in its dual, and it is equivalent - and even sometimes simpler - to study the dual instead.

¹The set of *Borel sets* of a topological space is the smallest set that is closed under complementation, countable union and contains all the open sets. In particular, it is closed under countable intersection and contains all the closed sets.

Lemma 3.3 below is an analogous non-topological characterisation of ‘topological circuits’ for graphs in general. We start explaining some topological concepts.

The *simplicial topology* of G is obtained from the disjoint union of copies $e \times (0, 1)$ of the unit interval, one for each edge e of G , by identifying two endpoints of these intervals if they correspond to the same vertex.

First we recall the definition of the *end space* $|G|$ from [13], and then we give an equivalent one using inverse limits. The point space of $|G|$ is the union of the set of vertex ends $\Omega(G)$, the vertex set $V(G)$ and a set $e \times (0, 1)$ for each edge e of G . A basis of this topology consists of the sets that are open in the simplicial topology of G (these do not contain vertex ends) together with sets of the form $C_{\vec{e}}(S, \omega)$, which are defined as follows: given a finite set of vertices S and a vertex end ω , by $C(S, \omega)$ we denote the component of $G - S$ in which ω lives. Let \vec{e} be a function from the set of those edges with exactly one endvertex in $C(S, \omega)$ to $(0, 1)$. The set $C_{\vec{e}}(S, \omega)$ consists of all vertices of $C(S, \omega)$, all vertex ends living in $C(S, \omega)$, the set $e \times (0, 1)$ for each edge e with both endvertices in $C(S, \omega)$, together with for each edge f with exactly one endvertex $t(f)$ in $C(S, \omega)$, the set of those points on $f \times (0, 1)$ with distance less than $\vec{e}(f)$ from $t(f)$. Note that $|G|$ is Hausdorff.

An *inverse system* of topological spaces consists of a partially ordered set (\mathcal{S}, \geq) of topological spaces together with continuous surjective functions $f: A \rightarrow B$ if $A \geq B$. Roughly, an *inverse limit* of such a system is a ‘smallest’ topological space T such that for every $A \in \mathcal{S}$ there is a continuous surjective function $\pi_A: T \rightarrow A$ such that these functions ‘commute’ with the functions of the inverse system. If any pair of elements of \mathcal{S} has an ‘upper bound’, then the inverse limit is uniquely determined by \mathcal{S} up to homeomorphism. The points of the inverse limit have the form $(x_A \mid A \in \mathcal{S})$, where $f(x_A) = x_B$ for all $A, B \in \mathcal{S}$ with $A \geq B$. Its topology is the coarsest so that all the projections π_A are continuous.

Given a finite vertex set W of G , by $G^+[W]$ we denote the (multi-) graph obtained from G by contracting all edges not incident with a vertex of W . Thus the vertex set of $G^+[W]$ is W together with the set of components of $G - W$. We consider $G^+[W]$ as a topological space endowed with the simplicial topology. If $U \subseteq W$, then by $f[W, U]$ we denote the surjective continuous map from $G^+[W]$ to $G^+[U]$, which is the identity on U , and maps each vertex w of $W \setminus U$ to the unique vertex component vertex of $G^+[U]$ whose component contains w . Moreover at interior points of edges e the map $f[W, U]$ commutes with the two fixed homeomorphisms to $(0, 1)$ of the two contraction graphs for e .

Theorem 3.1. $|G|$ is the inverse limit of the topological spaces $G^+[W]$ with respect to the maps $f[W, U]$.

Proof. For each vertex v of G , there is a point in the inverse limit which in the component for $G^+[W]$ takes the vertex whose branch set² contains v . This is the point corresponding to the vertex v . Similarly, there are points in the inverse limit corresponding to interior points of edges. All other points in the inverse limit correspond to havens of order $< \infty$ of G .³ As shown by Diestel and Kühn in [16, Theorem 2.2], these are precisely the vertex ends of G . Thus $|G|$ and the inverse limit have the same point set. It is straightforward to check that they carry the same topology. \square

In particular, $|G|$ has the following universal property: Suppose there is a topological space X and for each finite set W of vertices of G , a continuous function $f_W : X \rightarrow G^+[W]$ such that $f[W, U] \circ f_W = f_U$ for every $U \subseteq W$. Then there is a unique continuous function $f : X \rightarrow |G|$ such that $\pi_W \circ f = f_W$, where $\pi_W : |G| \rightarrow G^+[W]$ is the canonical projection.

A function f from S^1 to $|G|$ is *sparse* if $f^{-1}(v)$ never contains more than one point for each interior point v of an edge, and if there are two distinct points $x, y \in S^1$ with $f(x) = f(y)$, then there are two points z_1 and z_2 in different components of $S^1 - x - y$ both of whose f -values are different from $f(x)$ and not equal to interior points of edges.

Remark 3.2. (*Motivation*) Note that every injective function is sparse. However, because of the characterisation given in Lemma 3.3 below, it will turn out to be more convenient to work with sparse continuous functions instead of injective ones. The reason for this is a little technical: Roughly, the relation between sparse continuous functions and injective ones is similar to the relation between cuts and bonds. Indeed, the image of every sparse continuous function includes the image of an injective one. At some point in the argument later on it will be easier to work with cuts and sparse continuous functions in the first place and then use some ‘duality argument’ (Lemma 2.5, to be precise) to go from there to the minimal ones (meaning bonds and injective functions).

²If H is a minor of G , then the *branch set* of a vertex $v \in H$ is the (connected) set of those edges contracted onto v .

³A *haven* of order $k + 1$ in a graph G consists of a choice of a component of $G - S$ for every set S of at most k vertices. Moreover, it is required that any two such components share a vertex or *touch* (that is: there is an edge with one endvertex in the first component and one endvertex in the second component.)

Intuitively, the first property of a sparse function f says that f ‘traverses’ each edge only once, and the second says that it cannot ‘sit on a vertex’ for a nontrivial interval.

Let f from S^1 to $|G|$ be a sparse continuous function. Then f meets an edge e in an interior point if and only if it traverses this edge precisely once. The set of those edges e is called the *edge set of f* , denoted by $E(f)$. If f is a topological cycle, we call $E(f)$ a *topological circuit*. An edge set F is *geometrically connected* if $F \cap b \neq \emptyset$ for every finitely coverable cut b with the property that at least two components of $G - b$ contain edges of F . For example, if the closure of an edge set F in $|G|$ is connected in $|G|$, then F is geometrically connected.

Lemma 3.3. *A nonempty edge set X is the set of edges of a sparse continuous function f from S^1 to $|G|$ if and only if it meets every finitely coverable cut evenly and is geometrically connected.*

Proof. For the ‘only if’-implication, first note that the edge set of f is geometrically connected since connectedness is preserved under continuous images. Second, let F be a finitely coverable cut and let W be a finite vertex set covering it. If there is a sparse continuous function $f : S^1 \rightarrow |G|$, then $\pi_W \circ f : S^1 \rightarrow G^+[W]$ is also continuous. By continuity, $Y = E(\pi_W \circ f)$ has finite even degree at each vertex of W . Hence $Y \cap F$ is even. So $X \cap F = Y \cap F$ is even.

The ‘if’-implication is a consequence of Theorem 3.1: Suppose we have a geometrically connected set X meeting every finitely coverable cut evenly. Then for every finite vertex set W , the edge set $X \cap E(G^+[W])$ meets every cut of $G^+[W]$ evenly and is geometrically connected. Put another way, $X \cap E(G^+[W])$ is a finite connected set that has even degree at every vertex of $G^+[W]$. So it is eulerian, that is, $X \cap E(G^+[W])$ is the edge set of a sparse continuous function $f_W : S^1 \rightarrow G^+[W]$. We may assume that the f_W traverse every fixed edge with the same speed. Thus we can use a standard compactness argument to ensure that $f_U = f[W, U] \circ f_W$ for every $U \subseteq W$. Then the limit of the f_W is continuous by the above stated universal property of the limit and it is sparse by construction. \square

The simplest example of a finitely coverable cut is the set of edges incident with a fixed vertex. Thus the edge set of a sparse continuous function has even degree at each vertex by Lemma 3.3. Thus we get the following.

Corollary 3.4. *Given a sparse continuous function f , then for every finite vertex set W only finitely many components of $G - W$ contain vertices incident with edges of $E(f)$.*

Proof. Let X be the set of those edges of $E(f)$ incident with vertices of W . Note that X is finite by Lemma 3.3. If two components of $G - W$ contain vertices incident with edges of $E(f)$, then the incidence set s_D intersects X for every component D containing vertices incident with edges of $E(f)$, as $E(f)$ is geometrically connected by Lemma 3.3. Thus there are only finitely many such components D . \square

Having Lemma 3.3 and Corollary 3.4 in mind, the set F below can be thought of as the edge set of a topological cycle. Thus the following is an extension of the ‘Jumping arc’-Lemma [15]:

Lemma 3.5. *Let F be an edge set meeting every finitely coverable cut evenly such that for every finite vertex set W only finitely many components of $G - W$ contain vertices of $V(F)$. Let b be a cut which does not intersect F evenly. Then there is a vertex end in the closure of both F and b .*

Given a finite vertex set W and a component D of $G - W$, we denote by $v(D)$ the vertex of $G^+[W]$ with branch set D .

Proof. First we show that for every finite vertex set W there is a component D of $G - W$ such that the incidence set s_D contains infinitely many edges of both F and b . Suppose for a contradiction there is a vertex set W violating this. For a component D of $G - W$, let $X(D)$ be the set of those vertices in D incident with edges of b . Similarly, let $Y(D)$ be the set of those vertices in D incident with edges of F . Let U be the union of W with those $X(D)$ such that $Y(D)$ is infinite and those $Y(D)$ such that $Y(D)$ is finite.

By assumption $Y(D)$ is empty for all but finitely many D . Thus U is finite. For each component K of $G - U$ at least one of $X(K)$ and $Y(K)$ is empty. In particular, the graph $G^+[U]$ contains all edges of $b \cap F$.

Since F has even degree at each vertex of $G^+[U]$ and it has only finitely many edges in $G^+[U]$, it is a finite edge-disjoint union of cycles. Moreover, each of these cycles never uses vertices $v(K)$ for components K of $G - U$ with $X(K)$ nonempty. Hence each of these cycles extends to a cycle of G by only using additionally edges not in b . Since b is a cut of G , it meets each of these cycles evenly. Hence $b \cap F$ is even, which is the desired contradiction.

Hence for every finite vertex set W there is a component D_W of $G - W$ such that the incidence set s_{D_W} contains infinitely many edges of both F and b . Since for each set W there are only finitely many such components D_W , we can use a standard compactness argument to pick the components D_W with the additional property that if $U \subseteq W$, then $f[U, W](v(D_W)) = v(D_U)$. Thus the components D_W define a haven of order $< \infty$ of G , which defines

a vertex end ω , see [16, Theorem 2.2]. By construction the vertex end ω is in the closure of both F and b , completing the proof. \square

Lemma 3.6. *Let f be a sparse continuous function from S^1 to $|G|$ and let $x, y \in S^1$ such that $f(x)$ and $f(y)$ are distinct and not interior points of edges. Then for each connected component C of $S^1 - x - y$ there is an edge e_C of G such that $e_C \times (0, 1)$ is included in $f(C)$.*

Proof. We pick a finite vertex set W such that $\pi_W(f(x))$ is not equal to $\pi_W(f(y))$. Then the edge set of $\pi_W(f(C))$ is a finite walk in $G^+[W]$ from $\pi_W(f(x))$ to $\pi_W(f(y))$. Let e be an edge on that walk. Then $e \times (0, 1)$ is included in $f(C)$. \square

4 Proof of Theorem 1.4

Given a connected graph G , we fix a tree-decomposition $(T, (P_t \mid t \in V(T)))$ as in Theorem 2.1. For an undominated vertex end ω of G , we denote the unique end of T in which it lives by $\iota_T(\omega)$. It is straightforward to check that ι_T is a homeomorphism from the set of undominated vertex ends to $\Omega(T)$.

For each $t \in V(T)$, let $M(t)$ be the finite-cycle matroid of the torso H_t . Let $\bar{V}(t) = V_{M(t)}$ and $\bar{W}(t) = W_{M(t)}$. Thus $\bar{V}(t)$ consists of those finite edge sets of H_t that have even degree at every vertex, and $\bar{W}(t)$ consists of the cuts of H_t .

Remark 4.1. $\mathcal{T} = (T, \bar{V}, \bar{W})$ is a tree of binary finitary presentations. \square

The aim of this section is to prove Theorem 1.4 from the Introduction. For that we have to show for each Borel set Ψ of undominated vertex ends of G that certain sets C_Ψ and D_Ψ are the sets of circuits and cocircuits of a matroid. By Theorem 2.10, we know that $\Pi_{\iota_T(\Psi)}(\mathcal{T})$ presents some matroid. In this section we prove that the circuits and cocircuits of that matroid are given by C_Ψ and D_Ψ .

To build this bridge from $\Pi_{\iota_T(\Psi)}(\mathcal{T})$ to the sets C_Ψ and D_Ψ , we start as follows. We have the two topological spaces $\Omega(G)$ and $\Omega(T)$, which each have their own Borel sets. The next lemma shows that these two systems of Borel sets are compatible:

Lemma 4.2. *The set of dominated vertex ends of G is Borel. Furthermore, for any set Ψ of undominated vertex ends, Ψ is Borel in $\Omega(G)$ if and only if $\iota_T(\Psi)$ is Borel in $\Omega(T)$.*

To prove this lemma, we need some intermediate lemmas. By $B_k(s)$ we denote the ball of radius k around a fixed vertex s .

Lemma 4.3. *The graph $G[B_k(s)]$ has a spanning tree Y_k of diameter at most $2k$.*

Proof. Proving this by induction over k , we may assume that $G[B_{k-1}(s)]$ has a spanning tree Y_{k-1} of diameter at most $2k - 2$. Then Y_{k-1} together with all edges joining vertices in $B_k(s) \setminus B_{k-1}(s)$ to vertices in Y_{k-1} is a connected subgraph of $G[B_k(s)]$ with vertex set $B_k(s)$. Let Y_k be any of its spanning trees extending Y_{k-1} . Moreover, Y_k has diameter at most $2k$ by construction. \square

Lemma 4.4. *Let G be a graph with a fixed vertex s . The set Ω_k of those vertex ends dominated by some vertex in $B_k(s)$ is closed.*

Proof. In order to show that Ω_k is closed, we prove that its complement is open. For that it suffices to find for each ray R not dominated by some vertex in $B_k(s)$ some finite separator S_R disjoint from $B_k(s)$ that separates $B_k(s)$ from a tail of R .

Suppose for a contradiction that there is not such a finite separator S_R . Then we can recursively pick infinitely many $B_k(s)$ - R -paths that are vertex-disjoint except possibly their starting vertices. Let U be the set of their starting vertices. The set U is infinite because otherwise some $u \in U$ would dominate R , which is impossible. By Lemma 4.3, $G[B_k(s)]$ has a rayless spanning tree Y_k . Applying the Star-Comb-Lemma [15, Lemma 8.2.2] to Y_k and U , we find a vertex v in $G[B_k(s)]$ together with an infinite fan whose endvertices are in U . Enlarging this fan by infinitely many of the previously chosen $B_k(s)$ - R -paths, yields an infinite fan which witnesses that v dominates R , which is the desired contradiction. Thus there is such a finite set R_S for every ray R not dominated by some vertex in $B_k(s)$ and so Ω_k is closed. \square

Proof that Lemma 4.4 implies Lemma 4.2. By Lemma 4.4, the set of dominated vertex ends is a countable union of closed sets and thus Borel. Since the intersection of two Borel sets is Borel, a set of undominated vertex ends is Borel in $\Omega(G)$ if and only if it is Borel in the subspace topology of the undominated vertex ends inherited from $\Omega(G)$. So the ‘Furthermore’-part follows since Borelness is preserved by homeomorphisms. \square

The next step in our proof of Theorem 1.4 is to give a more combinatorial description of the set C_Ψ defined in the Introduction. For a set A , we denote

the set of minimal nonempty elements of A by A^{min} . Given a set Ψ of vertex ends of G , an edge set o is in \mathcal{C}_Ψ if o has only vertex ends of Ψ in its closure and it meets every finitely coverable cut evenly and it is geometrically connected. The next lemma implies that $\mathcal{C}_\Psi = \mathcal{C}_\Psi^{min}$.

Lemma 4.5. *Given a set Ψ of vertex ends of G , the following are equivalent for some nonempty edge set o .*

1. $o \in \mathcal{C}_\Psi$;
2. o is the edge set of a sparse continuous function from S^1 to $|G|$ that only has vertex ends from Ψ in the closure;
3. o is the edge set of a sparse continuous function from S^1 to $|G| \setminus \Psi^c$.

In particular, if o is minimal nonempty with one of these properties, then it is minimal nonempty with each of them. Furthermore o is minimal nonempty with one of these properties if and only if o is the edge set of a topological cycle in $|G| \setminus \Psi^c$.

Proof of Lemma 4.5. Clearly 2 and 3 are equivalent. And 1 and 2 are equivalent by Lemma 3.3. Thus 1,2 and 3 are equivalent.

To see the ‘Furthermore’-part, first note that the edge set of a topological cycle in $|G| \setminus \Psi^c$ is a minimal nonempty edge set satisfying 3. To see the converse, let o be a minimal edge set which is the edge set of a sparse continuous function f from S^1 to $|G| \setminus \Psi^c$. Suppose for a contradiction that f is not injective. Then there are two distinct points $x, y \in S^1$ with $f(x) = f(y)$. By sparseness of f , there are two points z_1 and z_2 in different components of $S^1 - x - y$ whose f -values are different from $f(x)$. By Lemma 3.6 applied first to x and z_1 and second to x and z_2 , for each of the two components C_1 and C_2 of $S^1 - x - y$ there is for each $i = 1, 2$ an edge e_i of G such that $e_i \times (0, 1)$ is included in $f(C_i)$.

We obtain the topological space K from $C_1 \cup \{x, y\} \subseteq S^1$ by identifying x and y . Note that K is homeomorphic to S^1 . Moreover, the restriction \bar{f} of f to $C_1 \cup \{x\}$ considered as a map from K to $|G|$ is continuous. However, the edge set of \bar{f} is included in the edge set of f without e_2 , violating the minimality of the edge set of f . Thus f is injective, and so o is the edge set of a topological cycle in $|G| \setminus \Psi^c$, completing the proof. \square

Let \mathcal{D}_Ψ be the set of cuts that do not have a vertex end of Ψ in their closure. Put another way, $d \in \mathcal{D}_\Psi$ if and only if d does not have a vertex

end of Ψ in its closure and it meets every finite cycle evenly. Note that $D_\Psi = \mathcal{D}_\Psi^{min}$. The next step in our proof of Theorem 1.4 is to relate \mathcal{C}_Ψ and \mathcal{D}_Ψ to the sets of $\iota_T(\Psi)$ -vectors of \mathcal{T} and $\iota_T(\Psi)^c$ -covectors of \mathcal{T} .

Lemma 4.6.

1. *The edge set of a finite cycle is an underlying vector of an \emptyset -pre-vector of \mathcal{T} ;*
2. *Any finitely coverable bond is an underlying covector of an \emptyset -pre-covector of \mathcal{T} .*

Proof. In this proof we use the tree order \leq on T induced by the root of T .

To see the second part, let d be a finitely coverable bond and let $V(G) = A \dot{\cup} B$ be a partition inducing d and let A' be a finite cover of d . Since G is connected, the partition is unique and both A and B are connected.

For $t \in V(T)$, let $x(t)$ be the set of crossing edges of the partition $V(P_t) = (A \cap V(P_t)) \dot{\cup} (B \cap V(P_t))$ in the torso H_t . Let S be the set of those nodes t such that A and B both meet $V(P_t)$.

Our aim is to show that (S, x) is an \emptyset -pre-covector of \mathcal{T} , which then by construction has underlying set d . By construction, $x(t) \in \overline{W}(t)$. It remains to verify the followings sublemmas.

Sublemma 4.7. *S is connected. Moreover, for each $st \in E(S)$, $x(s)$ contains an edge of the torso H_t .*

Sublemma 4.8. *S is rayless.*

Proof of Sublemma 4.7. It suffices to show for each $vw \in E(T)$ separating two vertices of S that $X = V(P_v) \cap V(P_w)$ contains vertices of both A and B . This follows from the fact that A and B are both connected and each has vertices in at least two components of $G - X$. \square

Proof of Sublemma 4.8. Suppose for a contradiction that S includes a ray $v_1 v_2 \dots$. By taking a subray if necessary we may assume that $v_i < v_{i+1}$. By the second property of strongly exhausting, the graph $G_i = (V[v_i], S[v_i v_{i+1}])$ is connected. Since A and B contain vertices of $P_{v_i} \subseteq G_i$, the graph G_i contains an edge of d . Then the finite cover A' of d has to contain a vertex of each G_i . Since A' is finite, there must a vertex in infinitely many G_i . However, this contradicts the first property of strongly exhausting. \square

To see the first part, let o be the edge set of a finite cycle. We shall define for each node $t \in V(T)$ an edge set $x(t)$, which plays a similar role as

in the last part. For that we need some preparation. Let $y(t) = o \cap E(P_t)$. Let $vw \in E(T)$ with $v < w$. Let J be the separation corresponding to the directed edge vw . Let $Z(vw)$ be the set of those vertices of o in $V(P_v) \cap V(P_w)$ such that of its two incident edges in o one is in J and the other in the complement of J . It can be shown that $|Z(vw)|$ is even; indeed just prove it by induction on the number of arcs consisting of edges in o that meet $V(P_v) \cap V(P_w)$ precisely in its two endvertices.

Recall that $H_v \cap H_w$ is a complete graph. So we can pick a perfect matching $M(vw)$ of $Z(vw)$ using only edges from $E(H_v) \cap E(H_w)$. We obtain $x(w)$ from $y(w)$ by adding all the sets $M(vw)$ where v is a neighbour of w . Let S be the set of those nodes w where $x(w)$ is nonempty.

Our aim is to show that (S, x) is an \emptyset -pre-vector of \mathcal{T} , which then by construction has underlying set o . Since o is finite, there are only finitely many nonempty $y(w)$. Since S is finite, it remains to verify the following sublemmas.

Sublemma 4.9. *$x(t)$ has even degree at each vertex of H_t .*

Sublemma 4.10. *S is connected.*

Proof of Sublemma 4.9. By the first property of strongly exhausting, $x(t)$ has even degree at all vertices v in $V(H_t) \cap V(H_s)$, where $st \in E(T)$ with $s < t$. Hence if t is maximal in S , then $x(t)$ has even degree at all vertices of H_t . Otherwise the statement follows inductively from the statement for all the upper neighbours. Indeed, let $v \in V(H_t) \setminus V(H_s)$, where $st \in E(T)$ with $s < t$. Then the degree of v in $x(t)$ is congruent modulo 2 to the degree of v in o plus the sum of the degrees of v in $x(u)$, where the sum ranges over all upper neighbours u of t . \square

Proof of Sublemma 4.10. It suffices to show for each $pq \in E(T)$ with $p < q$ separating two vertices of S that $x(p)$ contains a matching edge of $M(pq)$. Let X be the separation corresponding to the directed edge pq . Then the separator $\partial(X)$ of X separates two edges of o . Hence $\partial(X)$ must contain a vertex v_1 of o such that one of its two incident edges of o is in X and the other is in its complement. So $v_1 \in Z(pq)$ and so $x(p)$ must contain a matching edge of $M(pq)$. \square

\square

Corollary 4.11. *Every $\Psi^{\mathbb{C}}$ -covector d of \mathcal{T} is in \mathcal{D}_{Ψ} .*

Proof. First note that d has only vertex ends of $\Psi^{\mathcal{C}}$ in its closure. Since \mathcal{T} is tree of binary finitary presentations, we can apply Lemma 4.6 and Lemma 2.11. Thus d is a cut as it meets every finite cycle evenly. \square

Let \mathcal{F}_Ψ be the set of those edge sets o meeting every finitely coverable cut evenly such that for every finite vertex set W only finitely many components of $G - W$ contain vertices of $V(o)$. By Lemma 3.3 and Corollary 3.4, we have $\mathcal{C}_\Psi \subseteq \mathcal{F}_\Psi$.

Lemma 4.12. *Any nonempty $o \in \mathcal{F}_\Psi$ includes a nonempty element of \mathcal{C}_Ψ . Hence, $\mathcal{F}_\Psi^{\min} = \mathcal{C}_\Psi^{\min}$.*

Proof. We say that edges e and f of o are in the *same geometric component* if $o \cap d \neq \emptyset$ for every finitely coverable cut d such that e and f are in different components of $G - d$. It is straightforward to check that being in the same geometric component is an equivalence relation. Let u be the equivalence class of some element of o . It suffices to show that u is in \mathcal{C}_Ψ , which is implied by the following two sublemmas.

Sublemma 4.13. *u is meets every finitely coverable cut evenly.*

Sublemma 4.14. *u is geometrically connected.*

Before proving these two sublemmas, we give a construction that is used in the proof of both these sublemmas. Let $x \in u$ and let b be a finitely coverable cut. For all $z \in b \cap (o \setminus u)$, there is a finitely coverable cut b_z such that x and z are in different components of $G - b_z$. Let $V(G) = A \dot{\cup} B$ be a partition inducing b , and let $V(G) = A_z \dot{\cup} B_z$ be a partition inducing b_z such that x has both its endvertices in A_z . Let d be the cut consisting of those edges with precisely one endvertex in the intersection of A and the finitely many A_z . Note that d is finitely coverable. By construction $d \cap u = d \cap o$. Moreover, $b \cap u = d \cap u$ since any $y \in u$ has both its endvertices in any A_z .

Proof of Sublemma 4.13. Let b be a finitely coverable cut. Then $b \cap u = d \cap o$, and thus $b \cap u$ has even size. \square

Proof of Sublemma 4.14. Let b be a finitely coverable cut such that there are edges x and y of u in different components of $G - b$. Thus there is a partition $V(G) = A \dot{\cup} B$ inducing b such that x has both endvertices in A and y has both endvertices in B . Then x and y are in different components of $G - d$. As x and y are in the same geometric component, d meets o . Thus b meets u , completing the proof. \square

\square

Lemma 4.15. *Every Ψ -vector o of \mathcal{T} is in \mathcal{F}_Ψ .*

Proof. The set o is a finite symmetric difference of sets o_i , which are underlying sets of Ψ -pre-vectors (S_i, \bar{o}_i) . Note that S_i is locally finite as each \bar{o}_i is finite and for each $xy \in E(S_i)$, the set $\bar{o}_i(x)$ contains an edge of the torso of P_y . By the first property of strongly exhausting, this implies that o has finite degree at each vertex.

This implies that o meets every finitely coverable cut finitely; indeed, every edge in the intersection must have an endvertex from the finite cover and o has finite degree at any vertex. Since \mathcal{T} is a tree of binary finitary presentations, we can apply Lemma 4.6 and Lemma 2.11 to deduce that o meets every finitely coverable bond evenly. Since every cut is an edge-disjoint union of bonds, it follows that o meets every finitely coverable cut evenly.

It remains to show that for every o_i there is no finite vertex set W together with an infinite set \mathcal{A} of components of $G - W$ each containing a vertex of $V(o_i)$.

Suppose for a contradiction there is such a set W for o_i . Let Q be the smallest subtree of T containing the root r and all nodes q such that its part P_q contains a vertex of W . By the first property of strongly exhausting, Q is rayless. For each $A \in \mathcal{A}$, there is an edge z_A in $o_i \cap s_A$ ⁴. Let t_A be the unique node of T such that $z_A \in P_{t_A}$.

Since S_i is locally finite and Q is rayless, the forest $S_i \cap Q$ is finite. Since each set $\bar{o}_i(x)$ is finite, there can be only finitely many t_A that are in Q as all the z_A are distinct.

Hence we may assume that t_A is not in Q for all $A \in \mathcal{A}$ by taking a subset of \mathcal{A} if necessary. Let q_A be the last node on the unique t_A - Q -path and u_A be the node before that. By the second property of strongly exhausting, $(V[u_A], S[q_A u_A])$ is connected. Thus $V[u_A]$ is included in A . Thus the nodes u_A are distinct for different A . So the unique path from t_A to t_B for $A \neq B$ first goes to u_A , then it enters Q in q_A and then it further continues to t_B . As S_i is connected and $t_A, t_B \in S_i$, it must be that $q_A \in S_i$. So u_A is in S_i , as well. Since S_i is locally finite, the finite set $Q \cap S_i$ can have only finitely many neighbours in S_i . This contradicts the fact that infinitely many u_A are distinct. Hence o is in \mathcal{F}_Ψ . \square

Theorem 4.16. *Let Ψ be a Borel set of vertex ends of an infinite connected graph G that are all undominated. Then there is a matroid M whose set of circuits is \mathcal{C}_Ψ^{\min} and whose set of cocircuits is \mathcal{D}_Ψ^{\min} .*

⁴Recall that s_A denotes the set of edges incident with a vertex of A .

Proof. By Lemma 4.2, $\iota_T(\Psi)$ is Borel. Thus we apply Theorem 2.10 to the tree of presentations \mathcal{T} , yielding that $\Pi_{\iota_T(\Psi)}(\mathcal{T})$ presents a matroid M . Note that \mathcal{F}_Ψ and \mathcal{D}_Ψ satisfy (01) by Lemma 3.5. Hence by Corollary 4.11 and Lemma 4.15, we can apply Lemma 2.5 to \mathcal{F}_Ψ and \mathcal{D}_Ψ and M . As $\mathcal{F}_\Psi^{\min} = \mathcal{C}_\Psi^{\min}$ by Lemma 4.12, we get the desired result. \square

Proof of Theorem 1.4. First note that the set of undominated vertex ends is Borel by Lemma 4.2. By considering distinct connected components separately, we may assume that G is connected. By Lemma 4.5, \mathcal{C}_Ψ^{\min} is the set of topological cycles in $|G| \setminus \Psi^{\mathcal{C}}$. Thus Theorem 1.4 follows from Theorem 4.16. \square

Proof of Theorem 1.1. The first part of Theorem 1.1 is implied by Theorem 1.4. To see the ‘Moreover’-part, first note that the set of finite topological cycles of G together with the topological ends is equal to the set of edge sets of finite cycles of G . Hence from the topological ends matroid of G we can reconstruct the finite-cycle matroid of G . It is well-known that 3-connected graphs are isomorphic if and only if their finite-cycle matroids are isomorphic⁵, yielding the ‘Moreover’-part. \square

5 Consequences of Theorem 1.4

First, we prove for any graph G that the set of topological circuits is the set of circuits of a matroid if and only if G does not have a subdivision of the dominated ladder H . This theorem was already mentioned in the Introduction, see Corollary 1.3. We start with a couple of preliminary lemmas.

Lemma 5.1. *Let ω be a dominated vertex end of a graph G such that there are two vertex-disjoint rays R and S belonging to ω . Then G has a subdivision of H .*

Proof. Let v be a vertex dominating ω . By taking subrays if necessary, we may assume that v lies on neither R nor S . As R and S belong to the same vertex end, there are infinitely many vertex-disjoint paths P_1, P_2, \dots from R to S . We may assume that no P_i contains v . Let r_i be the endvertex of P_i on R and s_i be the endvertex of P_i on S . By taking a subsequence of the P_i if necessary, we can ensure that the order in which the r_i appear on R is

⁵This was observed by Thomassen [22] (just after the proof of Theorem 4.1).

r_1, r_2, \dots . Similarly, we may assume that the order in which the s_i appear on S is s_1, s_2, \dots .

Let Q_1, Q_2, \dots be an infinite fan from v to $R \cup S$. So for one of R or S , say R , there is an infinite fan Q'_1, Q'_2, \dots from v to it that avoids the other ray. As each P_i and each Q'_j is finite, we can inductively construct infinite sets $I, J \subseteq \mathbb{N}$ such that for $i \in I$ and $j \in J$ the paths P_i and Q'_j are vertex-disjoint.

Indeed, just consider the bipartite graph with left hand side $(P_i \mid i \in \mathbb{N})$ and right hand side $(Q'_j \mid j \in \mathbb{N})$ and put an edge between two paths P_i and Q'_j if they share a vertex. Now we use that each vertex of this bipartite graph has only finitely many neighbours on the other side to construct an independent set of vertices that intersects both sides infinitely. Indeed, for each finite independent set, there are two vertices, one on the left and one on the right, such that the independent set together with these two vertices is still independent. So there is such an infinite independent set and I is its set of vertices on the left and J is its set of vertices on the right.

Finally, v together with R, S and $(P_i \mid i \in I)$ and $(Q'_j \mid j \in J)$ give rise to a subdivision of H , which completes the proof. \square

Lemma 5.2. *Let o be a topological circuit that has the vertex end ω in its closure. Then there is a double ray both of whose tails belong to ω .*

This lemma already was proved in [5, Lemma 5.6] in a slightly more general context.

Proof of Corollary 1.3. If G has a subdivision of H , then as explained in the Introduction the set of topological circuits violates (C3).

Thus it remains to consider the case that G has no subdivision of H . Now we apply Theorem 1.4 with Ψ the set of undominated vertex ends, which is Borel by Lemma 4.2.

It suffices to show that every topological circuit o of G is a Ψ -circuit. So let ω be a vertex end in the closure of o . Then by Lemma 5.2 there is a double ray both of whose tails belong to ω . If ω was not in Ψ , then G would have a subdivision of H by Lemma 5.1. Thus ω is in Ψ . As ω was arbitrary, this shows that every vertex end in the closure of o is in Ψ . \square

Theorem 1.4 can also be used to extend a central result of [3] from countable graphs to graphs with a normal spanning tree as follows. Given a graph G with a normal spanning tree T , in [3] we constructed the Undomination graph $U = U(G, T)$. This graph has the pleasant property that it has few enough edges to have no dominated vertex end but enough edges to have

G as a minor. Moreover there is an inclusion \tilde{u} from the set of vertex ends of G to the set of vertex ends of U . By Theorem 1.4, for every Borel set Ψ , the Ψ -circuits of $U(G, T)$ are the circuits of a matroid. Now we use the following theorem.

Theorem 5.3 ([3, Theorem 9.9]). *Assume that $(U, \tilde{u}(\Psi))$ induces a matroid M . Then (G, Ψ) induces the matroid $M.E(G)$.*

We refer the reader to [3, Section 3] for a precise definition of when the pair (G, Ψ) consisting of a graph G and a vertex end set Ψ induces the matroid M . Very very roughly, this says that the set of certain ‘topological circuits’ which only use vertex ends from Ψ is the set of the circuits of M . However the topological space taken there is different from the one we take in this paper, so that the definition of topological circuit there does not match with the definition of topological circuit in this paper. For example, in this different notion a ray starting at a vertex v may also be a circuit if the vertex end it converges to is in Ψ and dominated by v . However these two notions of topological circuit are the same if no vertex dominates by a vertex end. Thus combining Theorem 5.3 and Theorem 1.4, we get the following.

Corollary 5.4. *Let G be a graph with a normal spanning tree and $\Psi \subseteq \Omega(G)$ such that $\tilde{u}(\Psi)$ is Borel. Then (G, Ψ) induces a matroid.*

For example, if we choose Ψ equal to the set of all vertex ends (dominated or not), then we get an interesting instance of this corollary. Like Theorem 1.4, this gives a recipe to associate a matroid (which we call $M_I(G)$) to every graph G that has a normal spanning tree which in general is neither finitary nor cofinitary. These two matroids need not be the same. For example, these two matroids differ for the graph obtained from the two-way infinite ladder by adding a vertex so that it dominates precisely one of the two vertex ends.

In fact the circuits of the matroid $M_I(G)$ can be described topologically, namely they are the edge sets of topological cycles in the topological space ITOP, see [13] for a definition of ITOP. About ITOP, we shall only need the following fact, which is not difficult to prove: Given a graph G , we denote by G_I , the multigraph obtained from G by identifying any two vertices dominating the same vertex end. It is not difficult to show that G and G_I have the same topological cycles. Thus in order to study when the topological cycles of G induce a matroid, it is enough to study this question for the graphs G_I . In what follows, we show that the underlying simple

subgraph G'_I of G_I always has a normal spanning tree. This will imply the following:

Corollary 5.5. *The topological circuits of ITOP induce a matroid for every graph.*

Indeed by Corollary 5.4, we just need to show that G'_I has a normal spanning tree. Let H' be the graph obtained from the dominated ladder H by adding a clone of the infinite degree vertex of H . Note that G'_I has no subdivision of H' . Thus G'_I has a normal spanning tree due to the following criterion:

Theorem 5.6 (Halin [19]). *If G is connected and does not have a subdivision of the complete graph on countably many vertices, then G has a normal spanning tree.*

6 Discussion

In this paper, we have shown that the topological circuits of any graph together with the topological ends form a matroid. This is one way how one can define the class of *graphic matroids*. Another one is the following:

Graph-like spaces were introduced by Thomassen and Vela [23], and further studied by Christian, Richter and Rooney [11, 12, 24, 25]. Graph-like spaces are topological spaces whose topological circuits very often form the set of circuits of a matroid, see [6] for details. These matroids are graphic in the sense that all their finite minors are cycle matroids of graphs. Moreover, all these matroids have to be *tame*.⁶ Conversely, any such matroid can be represented by a graph-like space:

Theorem 6.1 (Bowler, Carmesin, Christian). *A 3-connected matroid can be represented by a graph-like space if and only if it is tame and all its finite minors are cycle matroids of graphs.*

Since graphs together with the topological ends are examples of graph-like spaces, as well as all Psi-matroids of Theorem 1.4, the second approach deals with a larger class of matroids than the first. Having said this, it remains an open problem whether these two approaches lead to the same class of infinite matroids:

Open Question 6.2. *Is there a graph-like space inducing a 3-connected matroid which is not a minor of a Psi-matroid?*

Bowler showed that any such graph-like space cannot be compact [1].

⁶See [6] for an explanation why this is reasonable assumption in this context.

References

- [1] N. Bowler. Compact graph-like spaces. In preparation.
- [2] N. Bowler and J. Carmesin. An excluded minors method for infinite matroids. Preprint 2012, current version available at <http://arxiv.org/pdf/1212.3939v1>.
- [3] N. Bowler and J. Carmesin. Infinite matroids and determinacy of games. Preprint 2013, current version available at <http://arxiv.org/abs/1301.5980>.
- [4] N. Bowler and J. Carmesin. Infinite trees of matroids. Preprint 2014, available at <http://arxiv.org/pdf/1409.6627v1>.
- [5] N. Bowler and J. Carmesin. The ubiquity of psi-matroids. Preprint 2012, current version available at http://www.math.uni-hamburg.de/spag/dm/papers/ubiquity_psi_final.pdf.
- [6] N. Bowler, J. Carmesin, and R. Christian. Infinite graphic matroids. Preprint 2013, current version available at <http://arxiv.org/abs/1309.3735>.
- [7] H. Bruhn and R. Diestel. Duality in infinite graphs. *Comb., Probab. Comput.*, 15:75–90, 2006.
- [8] Henning Bruhn and Reinhard Diestel. Infinite matroids in graphs. *Discrete Math.*, 311(15):1461–1471, 2011.
- [9] Henning Bruhn, Reinhard Diestel, Matthias Kriesell, Rudi Pendavingh, and Paul Wollan. Axioms for infinite matroids. *Adv. Math.*, 239:18–46, 2013.
- [10] J. Carmesin. All graphs have tree-decompositions displaying their topological ends. Preprint 2014, available at <http://arxiv.org/pdf/1409.6640v4>.
- [11] R. Christian. *Infinite graphs, graph-like spaces and B-matroids*. PhD thesis, University of Waterloo, 2010.
- [12] Robin Christian, R. Bruce Richter, and Brendan Rooney. The planarity theorems of MacLane and Whitney for graph-like continua. *Electron. J. Combin.*, 17(1):Research Paper 12, 10, 2010.

- [13] R. Diestel. Locally finite graphs with ends: a topological approach. <http://arxiv.org/abs/0912.4213>.
- [14] R. Diestel. Spanning trees and k -connectedness. *J. Combin. Theory (Series B)*, 56:263–277, 1992.
- [15] R. Diestel. *Graph Theory* (4th edition). Springer-Verlag, 2010. Electronic edition available at: <http://diestel-graph-theory.com/index.html>.
- [16] R. Diestel and D. Kühn. Graph-theoretical versus topological ends of graphs. *J. Combin. Theory (Series B)*, 87:197–206, 2003.
- [17] R. Diestel and D. Kühn. On infinite cycles I. *Combinatorica*, 24:68–89, 2004.
- [18] A. Georgakopoulos. Infinite hamilton cycles in squares of locally finite graphs. *Advances in Mathematics*, 220:670–705, 2009.
- [19] R. Halin. Simplicial decompositions of infinite graphs. *Ann. Discrete Math.*, 3:93–109, 1978. Advances in graph theory (Cambridge Combinatorial Conf., Trinity Coll., Cambridge, 1977).
- [20] D.A. Higgs. Infinite graphs and matroids. Proceedings Third Waterloo Conference on Combinatorics, Academic Press, 1969, pp. 245 - 53.
- [21] M. Stein. Arboriticity and tree-packing in locally finite graphs. *J. Combin. Theory (Series B)*, 96:302–312, 2006.
- [22] C. Thomassen. Duality of infinite graphs. *J. Combin. Theory (Series B)*, 33:137–160, 1982.
- [23] C. Thomassen and A. Vella. Graph-like continua, augmenting arcs, and Menger’s theorem. *Combinatorica*, 29. DOI: 10.1007/s00493-008-2342-9.
- [24] A. Vella. *A fundamentally topological perspective on graph theory*. PhD thesis, University of Waterloo, 2005.
- [25] Antoine Vella and R. Bruce Richter. Cycle spaces in topological spaces. *J. Graph Theory*, 59(2):115–144, 2008.