The ubiquity of Psi-matroids

Nathan Bowler   Johannes Carmesin

May 13, 2013

Abstract

Solving (for tame matroids) a problem of Aigner-Horev, Diestel and Postle, we prove that every tame matroid M can be reconstructed from its canonical tree decomposition into 3-connected pieces, circuits and co-circuits together with information about which ends of the decomposition tree are used by M.

For every locally finite graph G, we show that every tame matroid whose circuits are topological circles of G and whose cocircuits are bonds of G is determined by the set Ψ of ends it uses, that is, it is a Ψ-matroid.

1 Introduction

There is a canonical matroid associated to any finite graph G, whose circuits are the (edge sets of) cycles in G and whose cocircuits are the bonds of G (the minimal nonempty cuts). However, when G is infinite there is no longer a canonical choice of matroid. Instead, there are at least two very natural choices. The first is the finite cycle matroid $M_{FC}(G)$, defined as above, which is always finitary (all its circuits are finite). The second is the topological cycle matroid $M_C(G)$ whose circuits are (edge sets of) topological circles in a suitable compactification $|G|$ of G and whose cocircuits are the finite bonds of G [13]. These topological circles arise naturally in the generalisations of some fundamental theorems from finite to infinite graphs [12]. The topological cycle matroid is always cofinitary (all its cocircuits are finite).

If G is an infinite planar graph then we cannot hope to find a dual graph G^* such that $M_{FC}(G*) = (M_{FC}(G))^*$ since the first of these is finitary whilst the second need not be. Instead, what happens (under some weak assumptions) is that there is a dual graph G^* such that $M_C(G^*) = (M_{FC}(G))^*$ and $M_{FC}(G^*) = (M_C(G))^*$ [8].

There is a whole range of matroids sitting in between these two. The new points added in the compactification $|G|$ are called the ends of G. For a set Ψ of ends and a countable graph G, we very often get a matroid $M_\Psi(G)$, called the Ψ-*matroid* of G, whose circuits are the topological circles in G that only use ends from Ψ. In particular, we always get a matroid in this way if Ψ is Borel [4]. The extreme cases – where Ψ is empty or the set of all ends – give rise to the finite and topological cycle matroids respectively. The Ψ-matroids also reflect
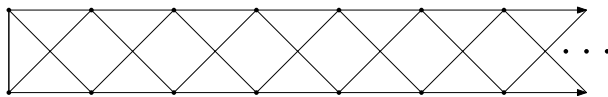
Figure 1: The graph Q

duality of graphs. If G^* is the dual of G then they have homeomorphic spaces of ends [10] and for any Borel set Ψ of ends we have $M_{\Psi\mathfrak{c}}(G^*) = (M_\Psi(G))^*$ [14, 4] (this specialises to the results above by taking Ψ to be empty or the set of all ends).

From now on we fix some locally finite graph G. All Ψ-matroids for G have the property that all their circuits are topological circles of G and all their cocircuits are bonds of G. We will call matroids with this property G-matroids. This paper is concerned with the question:

Question 1.1. *Characterise the G-matroids of locally finite graphs G.*

Not all G-matroids are Ψ-matroids.

Example 1.2. Let Q be the graph depicted in Figure 1. We say that two topological circuits of Q are *equivalent* if their symmetric difference is finite. Let $\mathcal{C}$ be any union of equivalence classes which includes all finite circuits. Then it can be shown that $\mathcal{C}$ is the set of circuits of a Q-matroid [7].

In cases like the above example, we can get a large set of badly behaved[1] matroids, each specified by a great deal of information 'at infinity'. In order to characterise the G-matroids of a general graph G we would require a specification of complex information distributed somehow over the ends of G, and this is currently intractable. However, there is a simple and natural restriction on the G-matroids which makes this problem much more tractable.

We say that a matroid M is *tame* if the intersection of any circuit with any cocircuit is finite. Otherwise, we say M is *wild*. The concept goes back to Dress [15], though not under this name: he introduced a different notion of infinite matroid to which tameness was fundamental. As recently as 2010 it was suggested that all infinite matroids might necessarily be tame [9]. Tameness and countability are the only extra restrictions we need to add to Minty's axiomatisation of finite matroids [16] to get an axiomatisation of the class of countable tame matroids: there is no need for an extra axiom requiring the existence of bases [4]. However, there are many wild matroids, as first shown in [5]. In fact, all but 2 of the Q-matroids in Example 1.2 are wild [7].

Nevertheless, the class of tame matroids is closed under duality and under taking minors, and includes most natural examples of matroids. For example, all finitary or cofinitary matroids and all Ψ-matroids are tame. Tameness gives a natural context for the consideration of representability of infinitary matroids.

---

[1] for example, such matroids can be non-binary in the sense that there are 3 circuits whose symmetric difference does not include a circuit

The notion of thin sums representability was introduced in [8] as a generalisation of representability to infinitary matroids. Although there is a wild matroid which is thin sums representable over $\mathbb{Q}$ but whose dual is not thin sums representable over any field [5], the class of tame thin sums matroids is closed under duality and under taking minors [1]. Moreover, forbidden minor characterisations extend readily from finite to tame thin sums matroids [3], so that for example a tame matroid is thin sums representable over $\mathbb{F}_2$ if and only if it does not have $U_{2,4}$ as a minor. Finally, tameness plays an essential role in the construction of an infinitary version of the class of graphic matroids [6], so that the badly behaved Q-matroids are not graphic in this sense.

We therefore restrict our attention to the class of tame G-matroids, for which we are now able to provide a simple characterisation.

Theorem 1.3. *The tame G-matroids are exactly the Ψ-matroids for G.*

However the set Ψ need not be Borel. The question of which sets Ψ of ends give rise to matroids is tied to subtle set theoretic questions about determinacy of games [4].

Restricting our attention to tame matroids also allows us to resolve a problem due to Aigner-Horev, Diestel and Postle from [2] about the reconstruction of connected matroids from their 3-connected pieces. Any finite connected matroid M can be decomposed canonically into a tree of pieces, each of which is 3-connected, a circuit or a cocircuit [11, 19]. Any two adjacent pieces share only a single edge, and M can be reconstructed from this tree by taking 2-sums along all these edges.

Recall that for any two matroids M_1 and M_2 that share only one edge, the *2-sum* $M_1 \oplus_2 M_2$ of M_1 and M_2 is the matroid whose edge set is the symmetric difference of those of M_1 and M_2 and whose circuits are of the following 3 types: circuits of M_1 avoiding e, circuits of M_2 avoiding e, and symmetric differences of an M_1-circuit containing e with an M_2-circuit containing e. This construction is associative, in the sense that if M_1 and M_2 meet in only one edge and M_2 and M_3 meet in only one edge, and M_1 has no edge in common with M_3, then $(M_1 \oplus_2 M_2) \oplus_2 M_3 = M_1 \oplus_2 (M_2 \oplus_2 M_3)$. Because of this associativity, it doesn't matter in what order we take the 2-sums at the edges of the tree: we always get back the original matroid M.

Aigner-Horev, Diestel and Postle partially extended this result to infinite matroids: they were able to show that there is such a canonical tree decomposition of any connected matroid [2]. It is a little surprising that the structure obtained is a genuine graph-theoretic tree, rather than one of the more order-theoretic or topological notions of infinite tree discussed in [17]. However, reconstruction of the original matroid from this tree is not so straightforward if the tree is infinite. For example, every Q-matroid decomposes into a ray of pieces, each of which is isomorphic to $M(K_4)$, as in Figure 2.

This example shows that the tree decomposition alone does not provide enough information to reconstruct the matroid: more information is needed. In [7], we answer the question of which extra information is needed to carry out this reconstruction. The answer is complicated and beyond the scope of this
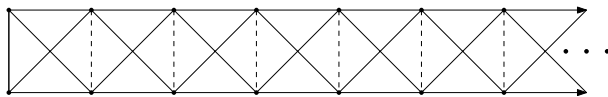
Figure 2: The tree decomposition of the graph Q

paper. However, if we once more restrict our attention to tame matroids then there is a much more natural solution. In this paper, we give a self-contained account of this more natural solution, for which the necessary arguments are much simpler than in [7].

The tree alone is still not enough information - the finite and topological cycle matroids of Q are both tame, and they give rise to the same tree of matroids. Just as in the graphic case, we may think of the topological cycle matroid as the matroid we get by allowing the end of the ray to be used by circuits and the finite cycle matroid as the matroid we get by forbidding circuits to use the end. This suggests what is in fact the right answer: the extra information we need is simply the set Ψ of ends of the tree which may be used by circuits.

More precisely, in Section 3 we give a construction which can be thought of as taking infinitely many 2-sums simultaneously. Given a suitable tree $\mathcal{T}$ of matroids and a suitable set Ψ of ends of $\mathcal{T}$, we showed in [4] that this construction allows us to build a matroid $M_\Psi(\mathcal{T})$ by (roughly speaking) sticking together the matroids along the edges of the tree and only allowing circuits to use the ends in Ψ. We can show that this construction suffices to rebuild any tame matroid from its canonical decomposition into circuits, cocircuits, and 3-connected pieces, together with information about which ends are used by circuits:

Theorem 1.4. *Let N be a tame matroid and let $\mathcal{T}$ be the tree of matroids $\mathcal{T}$ arising from the canonical tree decomposition of N.*

Then there is some $\Psi \subseteq \Omega(\mathcal{T})$ such that $N = M_\Psi(\mathcal{T})$.

The proof that the Ψ-matroids of a locally finite graph really are matroids also relied on gluing together an infinite tree of finite pieces. In this case, the tree structure arose from a tree decomposition of the graph. So both of the results mentioned above say that, for some particular tree structure, if we have any tame matroid whose circuits and cocircuits all fit, in some sense, with that tree structure, then this constrains the matroid to be of a very special type, which we call a Ψ-matroid. We give a general result of this type for trees of matroids in which any two adjacent matroids share at most one edge. See Theorem 4.10.

The paper is arranged as follows: we begin by recalling some preliminary results about infinite matroids in general and Ψ-matroids in particular in Section 2. Then we recall the basic theory of trees of matroids, including the infinitary version of 2-sums, in Section 3. We prove Theorem 1.4 in Section 4 and Theorem 1.3 in Section 5. We conclude by discussing open questions and the future outlook in Section 6.

2 Preliminaries

Throughout, notation and terminology for (infinite) graphs are those of [13], and for matroids those of [18, 9].

2.1 Infinite matroids

M always denotes a matroid and $E(M)$ (or just E), $\mathcal{I}(M)$ and $\mathcal{C}(M)$ denote its ground set and its sets of independent sets and circuits, respectively.

A set system $\mathcal{I} \subseteq \mathcal{P}(E)$ is the set of independent sets of a matroid if and only if it satisfies the following *independence axioms* [9].

(I1) $\emptyset \in \mathcal{I}(M)$.

(I2) $\mathcal{I}(M)$ is closed under taking subsets.

(I3) Whenever $I, I' \in \mathcal{I}(M)$ with I' maximal and I not maximal, there exists an $x \in I' \setminus I$ such that $I + x \in \mathcal{I}(M)$.

(IM) Whenever $I \subseteq X \subseteq E$ and $I \in \mathcal{I}(M)$, the set $\{I' \in \mathcal{I}(M) \mid I \subseteq I' \subseteq X\}$ has a maximal element.

A set system $\mathcal{C} \subseteq \mathcal{P}(E)$ is the set of circuits of a matroid if and only if it satisfies the following *circuit axioms* [9].

(C1) $\emptyset \notin \mathcal{C}$.

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

(C3) (Circuit elimination) Whenever $X \subseteq o \in \mathcal{C}(M)$ and $\{o_x \mid x \in X\} \subseteq \mathcal{C}(M)$ satisfies $x \in o_y \Leftrightarrow x = y$ for all $x, y \in X$, then for every $z \in o \setminus \left(\bigcup_{x \in X} o_x \right)$ there exists a $o' \in \mathcal{C}(M)$ such that $z \in o' \subseteq \left(o \cup \bigcup_{x \in X} o_x \right) \setminus X$.

(CM) $\mathcal{I}$ satisfies (IM), where $\mathcal{I}$ is the set of those subsets of E not including an element of $\mathcal{C}$.

Lemma 2.1. *Let M be a matroid and s be a base. Let o_e and b_f a fundamental circuit and a fundamental cocircuit with respect to s, then*

   1. $o_e \cap b_f$ is empty or $o_e \cap b_f = \{e, f\}$ and

   2. $f \in o_e$ if and only if $e \in b_f$.

Proof. To see the first note that $o_e \subseteq s + e$ and $b_f \subseteq (E \setminus s) + f$. So $o_e \cap b_f \subseteq \{e, f\}$. As a circuit and a cocircuit can never meet in only one edge, the assertion follows.

To see the second, first let $f \in o_e$. Then $f \in o_e \cap b_f$, so by (1) $o_e \cap b_f = \{e, f\}$ and so $e \in b_f$. The converse implication is the dual statement of the above implication. $\qquad\square$

Lemma 2.2. *For any circuit o containing two edges e and f, there is a cocircuit b such that $o \cap b = \{e, f\}$.*

Proof. As $o - e$ is independent, there is a base including $o - e$. By Lemma 2.1, the fundamental cocircuit of f of this base intersects o in e and f, as desired. $\square$

Lemma 2.3. *Let M be a matroid with ground set $E = C \dot\cup X \dot\cup D$ and let o' be a circuit of $M' = M/C \backslash D$. Then there is an M-circuit o with $o' \subseteq o \subseteq o' \cup C$.*

Proof. Let s be any M-base of C. Then $s \cup o'$ is M-dependent since o' is M'-dependent. On the other hand, $s \cup o' - e$ is M-independent whenever $e \in o'$ since $o' - e$ is M'-independent. Putting this together yields that $s \cup o'$ contains an M-circuit o, and this circuit must not avoid any $e \in o'$, as desired. $\square$

A *scrawl* is a union of circuits. In [3], (infinite) matroids are axiomatised in terms of scrawls. The set $\mathcal{S}(M)$ denotes the set of scrawls of the matroid M. Dually a *coscrawl* is a union of cocircuits. Since no circuit and cocircuit can meet in only one element, no scrawl and coscrawl can meet in only one element. In fact, this property gives us a simple characterisation of scrawls in terms of coscrawls and vice versa.

Lemma 2.4. *[3] Let M be a matroid, and let $w \subseteq E$. The following are equivalent:*

1. *w is a scrawl of M.*

2. *w never meets a cocircuit of M just once.*

3. *w never meets a coscrawl of M just once.*

Proof. It is clear that (1) implies (3) and (3) implies (2), so it suffices to show that (2) implies (1). Suppose that (2) holds and let $e \in w$. Then in the minor $M/(w-e) \backslash (E \backslash w)$ on the groundset $\{e\}$, e cannot be a co-loop, by the dual of Lemma 2.3 and (2). So e must be a loop, and by Lemma 2.3 there is a circuit o_e with $e \in o_e \subseteq w$. Thus w is the union of the o_e, and so is a scrawl.
$\square$

Lemma 2.5. *Let M be a matroid and $\mathcal{C}, \mathcal{D} \subseteq \mathcal{P}(E)$ such that every M-circuit is a union of elements of $\mathcal{C}$, every M-cocircuit is a union of elements of $\mathcal{D}$ and $|C \cap D| \neq 1$ for every $C \in \mathcal{C}$ and every $D \in \mathcal{D}$.*
Then $\mathcal{C}(M) \subseteq \mathcal{C} \subseteq \mathcal{S}(M)$ and $\mathcal{C}(M^*) \subseteq \mathcal{D} \subseteq \mathcal{S}(M^*)$

Proof. We begin by showing that $\mathcal{C}(M) \subseteq \mathcal{C}$. For any circuit o of M, pick an element e of o. Since o is a union of elements of $\mathcal{C}$ there is $o' \in \mathcal{C}$ with $e \in o' \subseteq o$. Suppose for a contradiction that o' isn't the whole of o, so that there is $f \in o \backslash o'$. By Lemma 2.2 there is some cocircuit b of M with $o' \cap b = \{e\}$. Then we can find $b' \in \mathcal{D}$ with $e \in b' \subseteq b$, and so $o' \cap b' = \{e\}$, giving the desired contradiction. Similarly we obtain that $\mathcal{C}(M^*) \subseteq \mathcal{D}$.

The fact that $\mathcal{C} \subseteq \mathcal{S}(M)$ is immediate from Lemma 2.4 since $\mathcal{C}(M^*) \subseteq \mathcal{D}$, and the proof that $\mathcal{D} \subseteq \mathcal{S}(M^*)$ is similar. $\square$

2.2 Ψ-matroids

In this subsection we shall review the definitions of Ψ-circuits and $\Psi^{\complement}$-bonds for a locally finite graph G with a specified set Ψ of ends. Much of what we say will be a review of the early parts of [14] and [4]. We say that two rays in G are *equivalent* if they cannot be separated by removing finitely many vertices from G. In other words, two rays are equivalent if they may be joined by infinitely many vertex-disjoint paths. An *end* of G is an equivalence class of rays under this relation, and the set of ends of G is denoted $\Omega(G)$.

Let d be the distance function on $V(G) \sqcup (0,1) \times E(G)$ considered as the ground set of the simplicial 1-complex formed from the vertices and edges of G. We define a topology on the set $V(G) \sqcup \Omega(G) \sqcup (0,1) \times E(G)$ by taking basic open neighbourhoods as follows:

- For $v \in V(G)$, the basic open neighbourhoods of v are the ϵ-balls $B_\epsilon(v) = \{x | d(v,x) < \epsilon\}$ for $\epsilon \leq 1$.

- For $(x,e) \in (0,1) \times E$ we say (x,e) is an *interior point* of e, and take the basic open neighbourhoods to be the ϵ-balls about (x,e) with $\epsilon \leq \min(x, 1-x)$.

- For $\omega \in \Omega(G)$, the basic open neighbourhoods of ω will be parametrised by the finite subsets S of $V(G)$. Given such a subset, we let $C(S, \omega)$ be the unique component of $G - S$ that contains a ray from ω, and let $\hat{C}(S, \omega)$ be the set of all vertices and inner points of edges contained in or incident with $C(S, \omega)$, and of all ends represented by a ray in $C(S, \omega)$. We take the basic open neighbourhoods of ω to be the sets $\hat{C}(S, \omega)$.

We call the topological space obtained in this way $|G|$.

For any set Ψ of ends of G, we set $\Psi^{\complement} = \Omega(G) \setminus \Psi$ and $|G|_\Psi = |G| \setminus \Psi^{\complement}$. Since G is locally finite, $|G|_\Psi$ can be given the structure of a graph-like space in the sense of [6] (and closely related to the earlier work of [20]): see [4].

A Ψ-*circuit* of G is an edge set whose $|G|_\Psi$-closure is homeomorphic to the unit circle. A $\Psi^{\complement}$-*bond* of G is a bond of G that has no ends from Ψ in its closure.

We say that (G, Ψ) *induces a matroid* M if $E(M) = E(G)$ and the M-circuits are the Ψ-circuits and the M-cocircuits are the $\Psi^{\complement}$-bonds. In this case, we call M the Ψ-*matroid* $M_\Psi(G)$ of G. In the following sense the Ψ-circuits and $\Psi^{\complement}$-bonds always behave like the circuits and cocircuits of a matroid.

Lemma 2.6. *No Ψ-circuit meets any $\Psi^{\complement}$-bond in a single edge.*

Proof. Suppose for a contradiction that some Ψ-circuit o meets some $\Psi^{\complement}$-bond b in a single edge f

Then $|G|_\Psi$ with all the interior points of edges of b removed has two connected components, namely the two sides of the bond. This contradicts the fact that $o - f$ is connected and contains both endvertices of f. $\square$

More is shown in [4]. If Ψ is Borel, then the Ψ-matroid $M_\Psi(G)$ always exists.

3 Trees of matroids

In this section we review the relationship between trees of matroids and tree decompositions of matroids.

Definition 3.1. A *tree $\mathcal{T}$ of matroids* consists of a tree T, together with a function M assigning to each node t of T a matroid $M(t)$ on 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 *dummy edges* of $M(t)$: thus the set of such dummy edges is $\bigcup_{tt' \in E(T)} E(tt')$.

The idea is that the dummy edges are to be used only to give information about how the matroids are to be pasted together, but they will not be present in the final pasted matroid, which will have ground set $E(\mathcal{T})$.

Definition 3.2. If T is a tree, and tu is a (directed) edge of T, we take $T_{t \to u}$ to be the connected component of $T - t$ that contains u. If $\mathcal{T} = (T, M)$ is a tree of matroids, we take $\mathcal{T}_{t \to u}$ to be the tree of matroids $(T_{t \to u}, M \!\restriction_{T_{t \to u}})$.

Definition 3.3. A tree $\mathcal{T} = (T, M)$ of matroids is *of overlap 1* if, for every edge tt' of T, $|E(tt')| = 1$. In this case, we denote the unique element of $E(tt')$ by $e(tt')$.

Given a tree of matroids of overlap 1 as above, a *precircuit* of $\mathcal{T}$ consists of a connected subtree C of T together with a function o assigning to each vertex t of C a circuit of $M(t)$, such that for any vertex t of C and any vertex t' adjacent to t in T, $e(tt') \in o(t)$ if and only if $t' \in C$. Given a set Ψ of ends of T, such a precircuit is called a Ψ-*precircuit* if all ends of C are in Ψ. The set of Ψ-precircuits is denoted $\overline{\mathcal{C}}(\mathcal{T}, \Psi)$.

Any Ψ-precircuit (C, o) has an *underlying set* $\underline{(C, o)} = E \cap \bigcup_{t \in V(C)} o(t)$. Minimal nonempty subsets of E arising in this way are called Ψ-*circuits* of $\mathcal{T}$. The set of Ψ-circuits of $\mathcal{T}$ is denoted $\mathcal{C}(\mathcal{T}, \Psi)$.

Definition 3.4. Let $\mathcal{T} = (T, M)$ be a tree of matroids. Then the *dual* $\mathcal{T}^*$ of $\mathcal{T}$ is given by (T, M^*), where M^* is the function sending t to $(M(t))^*$. For a subset C of the ground set, the tree of matroids $\mathcal{T}/C$ obtained from $\mathcal{T}$ by *contracting C* is given by $(T, M/C)$, where M/C is the function sending t to $M(t)/(C \cap E(t))$. For a subset D of the ground set, the tree of matroids $\mathcal{T} \backslash D$ obtained from $\mathcal{T}$ by *deleting D* is given by $(T, M \backslash D)$, where $M \backslash D$ is the function sending t to $M(t) \backslash (C \cap E(t))$.

Lemma 3.5. *For any tree $\mathcal{T}$ of matroids, $\mathcal{T} = \mathcal{T}^{**}$. For any disjoint subsets C and D of the ground set of $\mathcal{T}$ we have $(\mathcal{T}/C)^* = \mathcal{T}^* \backslash C$, $(\mathcal{T} \backslash D)^* = \mathcal{T}^*/D$ and $\mathcal{T}/C \backslash D = \mathcal{T} \backslash D/C$. If $\mathcal{T}$ has overlap 1 and $(\mathcal{T}, \Psi)$ induces a matroid M, then $(\mathcal{T}/C \backslash D, \Psi)$ induces the matroid $M/C \backslash D$ and $(\mathcal{T}^*, \Psi^{\complement})$ induces the matroid M^*.* $\qquad\square$

We will sometimes use the expression $\Psi^{\complement}$-*cocircuits of* $\mathcal{T}$ for the $\Psi^{\complement}$-circuits of $\mathcal{T}^*$. If there is a matroid whose circuits are the Ψ-circuits of $\mathcal{T}$ and whose cocircuits are the $\Psi^{\complement}$-cocircuits of $\mathcal{T}$ then we will call that matroid the Ψ-*matroid* for $\mathcal{T}$, and denote it $M_\Psi(\mathcal{T})$.

Lemma 3.6 (Lemma 5.5, [4])**.** *Let* $\mathcal{T} = (T, M)$ *be a tree of matroids of overlap 1,* Ψ *a set of ends of* T*, and let* (C, o) *and* (D, b) *be respectively a* Ψ-*precircuit of* $\mathcal{T}$ *and a* $\Psi^{\complement}$-*precircuit of* $\mathcal{T}^*$*. Then* $|\underline{(C, o)} \cap \underline{(D, b)}| \neq 1$.

Theorem 3.7 ([4])**.** *If* $\mathcal{T} = (T, M)$ *is a tree of matroids of overlap 1, and* Ψ *is a Borel set of ends of* T*, then there is a* Ψ-*matroid for* $\mathcal{T}$*.*

Remark 3.8. *In particular, we always get a* Ψ-*matroid when* Ψ *is the empty set or the set* $\Omega(T)$ *of all ends of* T*. It is clear that if each matroid* $M(t)$ *is finitary, then so is* $M_\emptyset(\mathcal{T})$*.*

So far we have discussed how to construct matroids by gluing together trees of 'smaller' matroids. Now we turn to a notion, taken from [2], of a decomposition of a matroid into a tree of such smaller parts.

Definition 3.9. A *tree decomposition of adhesion 2* of a matroid N consists of a tree T and a partition $R = (R_v)_{v \in V(T)}$ of the ground set E of N such that for any edge tt' of T the partition $(\bigcup_{v \in V(T_{t \to t'})} R_v, \bigcup_{v \in V(T_{t' \to t})} R_v)$ is a 2-separation of N.

Given such a tree decomposition, and a vertex v of T, we define a matroid $M(v)$, called the *torso* of T at v, as follows: the ground set of $M(v)$ consists of R_v together with a new edge $e(vv')$ for each edge vv' of T incident with v. For any circuit o of N not included in any set $\bigcup_{t \in V(T_{v \to v'})} R_v$, we have a circuit $\hat{o}(v)$ of $M(v)$ given by $(o \cap R_v) \cup \{e(vv') \in E(v) | o \cap \bigcup_{t \in V(T_{v \to v'})} R_v \neq \emptyset\}$. These are all the circuits of $M(v)$.

In this way we get a tree of matroids $\mathcal{T}(N, T, R) = (T, v \mapsto M(v))$ of overlap 1 from any tree decomposition of adhesion 2. For any circuit o of N we get a corresponding precircuit $(S_o, \hat{o})$, where S_o is just the subtree of T consisting of those vertices v for which $\hat{o}(v)$ is defined.

Note that $\underline{(S_o, \hat{o})} = o$. Each $M(v)$ really is a matroid [2, §4, §8], isomorphic to a minor of N [7], and that $\mathcal{T}(N^*, T, R) = (\mathcal{T}(N, T, R))^*$. [2] also contains the following theorem.

Theorem 3.10 (Aigner-Horev, Diestel, Postle)**.** *For any matroid* N *there is a tree decomposition* $\mathcal{D}(N)$ *of adhesion 2 of* N *such that all torsos have size at least 3 and are either circuits, cocircuits or 3-connected, and in which no two circuits and no two cocircuits are adjacent in the tree. This decomposition is unique in the sense that any other tree decomposition with these properties must be isomorphic to it.*

The above theorem is a generalisation to infinite matroids of a standard result about finite matroids [11, 19]. If N is a finite matroid, it is possible to reconstruct N from the decomposition $\mathcal{D}(N)$. However, as noted in the

introduction, it is not in general possible to reconstruct N from $\mathcal{D}(N)$ if N is infinite. Our aim in the next section will be to show that if N is tame, then not much extra information is needed to recover N. All we need is the set Ψ consisting of those ends of T that appear in the closure of some circuit of N.

4   Reconstruction

Let N be a tame matroid and let (T, R) be a tree decomposition of N of adhesion 2. We begin by considering the case that T is a ray $t_1, t_2, \ldots$. In this case, we can show that the tree $\mathcal{T}(N, T, R)$ is well behaved.

Definition 4.1. A precircuit (S, o) for a tree $\mathcal{T} = (T, M)$ of matroids of overlap 1 is called a *phantom precircuit* if there is an edge tt' of S such that $o(v) \cap E(\mathcal{T}) = \emptyset$ for $v \in V(S_{t \to t'})$.

$\quad$ $\mathcal{T} = (T, M)$ is *nice* if neither $\mathcal{T}$ nor $\mathcal{T}^*$ has any phantom precircuits.

$\quad$ Note that $\mathcal{T} = (T, M)$ is nice iff there is not $tt' \in E(T)$ such that in $\mathcal{T}_{t \to t'} = (T_{t \to t'}, M\restriction_{V(T_{t \to t'})})$ the edge $e(tt')$ is either a loop in $M_{\Omega(\mathcal{T}_{t \to t'})}(\mathcal{T}_{t \to t'})$ or a coloop $M_{\emptyset}(\mathcal{T}_{t \to t'})$.

Lemma 4.2. *Let N be a matroid with a tree decomposition (T, R) of adhesion 2.*

1. *For every N-circuit o its corresponding precircuit $(S_o, \hat{o})$ is not phantom.*

2. *If T is a ray, and there is a circuit o and a cocircuit b of N that both have edges in infinitely many of the R_v, then $\mathcal{T}(N, T, R)$ is nice.*

Proof. (1) follows from the definition of S_o.

$\quad$ For (2), let $T = t_1, t_2, \ldots$ be a ray. Now suppose for a contradiction that there is a phantom precircuit (S_c, c). Then for all sufficiently large n, the circuit $c(t_n)$ consists of $e(t_{n-1}t_n)$ and $e(t_n t_{n+1})$. In other words, $e(t_{n-1}t_n)$ and $e(t_n t_{n+1})$ are in parallel.

$\quad$ So $c(t_n) \subseteq \hat{o}(t_n)$, hence $c(t_n) = \hat{o}(t_n)$. This contradicts (1). The case that there is a phantom precocircuit (S_c, c) is similar. Hence $\mathcal{T}(N, T, R)$ is nice. $\quad\square$

Lemma 4.3. *Let $\mathcal{T} = (T, M)$ be a nice tree of matroids, then every $\emptyset$-circuit is an $\Omega(T)$-circuit.* $\quad\square$

$\quad$ By duality, an analogue of Lemma 4.3 is also true for cocircuits.

Lemma 4.4. *Let $\mathcal{T} = (T, M)$ be a nice tree of matroids, and N be a matroid such that $\mathcal{C}(N) \subseteq \mathcal{C}(M_{\Omega(T)}(\mathcal{T}))$ and $\mathcal{C}(N^*) \subseteq \mathcal{C}(M_{\emptyset}^*(\mathcal{T}))$.*
$\quad$ *Then $\mathcal{C}(M_{\emptyset}(\mathcal{T})) \subseteq \mathcal{C}(N)$ and $\mathcal{C}(M_{\Omega(T)}^*(\mathcal{T})) \subseteq \mathcal{C}(N^*)$.*

Proof. By duality, it suffices to prove only that $\mathcal{C}(M_{\emptyset}(\mathcal{T})) \subseteq \mathcal{C}(N)$. So let $o \in \mathcal{C}(M_{\emptyset}(\mathcal{T}))$. Since o never meets an element of $\mathcal{C}(M_{\emptyset}^*(\mathcal{T}))$ just once, it never meets an element of $\mathcal{C}(N^*)$ just once. Hence o includes an N-circuit o' by Lemma 2.4. Thus $o' \in \mathcal{C}(M_{\Omega(T)}(\mathcal{T}))$. By Lemma 4.3, we must have $o' = o$. So $o \in \mathcal{C}(N)$, as desired. $\quad\square$

Lemma 4.5. *Let N be a tame matroid with a tree decomposition (T, R) of adhesion 2. Assume that $T = t_1 t_2 \ldots$ is a ray.*

Then there are not a circuit o and a cocircuit b of N that both converge to the end ω of T.

Indeed, either $N = M_\emptyset(\mathcal{T}(N, T, R))$ or $N = M_{\{\omega\}}(\mathcal{T}(N, T, R))$

Proof. Suppose for a contradiction that there are such o and b. Then there are $l < m < n$ and $e_l, e_m, e_n \in E(N)$ such that $e_l \in b \cap E(t_l)$, and $e_m \in o \cap E(t_m)$, and $e_n \in b \cap E(t_n)$. Using the tameness of N, we make these choices in such a way that for any $i \geq m$, the intersection of $o \cap b$ with $E(t_i)$ is empty. We may also assume that o has an edge in some $E(t_k)$ with $k < m$, so that both dummy edges of $E(t_m)$ are in $\hat{o}(t_m)$.

Now b is a cocircuit of $M(\mathcal{T}(N, T, R), \emptyset)$ since $(S_b, \hat{b})$ is a precocircuit, and there cannot be a precocircuit whose cocircuit at any node t is a subset of $\hat{b}(t)$. By the dual of Lemma 2.2 there is some $M_\emptyset(\mathcal{T})$-circuit o_b meeting b only in e_l and e_n. Note that o_b is also a circuit of N by Lemma 4.4 since $\mathcal{T}(N, T, R)$ is nice by Lemma 4.2.

Now we build an $\emptyset$-precircuit $(S_C, \hat{C})$ as follows. First we set $S_C = (S_{o_b} \setminus \{t_1 \ldots t_m\}) \cup (S_o \cap \{t_1 \ldots t_m\})$. We take $\hat{C}(t_j) = \hat{o}_b(t_j)$ for $j > m$, and $\hat{C}(t_j) = \hat{o}(t_j)$ for $j \leq m$. Let C be the underlying circuit of $(S_C, \hat{C})$. Note that C is a circuit of $M_\emptyset(\mathcal{T})$ and so also a circuit of N by Lemma 4.4 and Lemma 4.2 as before.

We now apply circuit elimination in N to the circuits o and C, eliminating the edge e_m and keeping the edge e_n. Call the resulting circuit C'.

If $t_m \in S_{C'}$, then $\hat{C}'(t_m) \subseteq \hat{o}(t_m) - e_m$ (since both dummy edges of $E(t_m)$ are in $\hat{o}(t_m)$), which is impossible. So $S_{C'} \subseteq \{t_{m+1}, t_{m+2}, \ldots\}$. Hence $C' \cap b = \{e_n\}$, which is also impossible.

We have now established that there cannot be a circuit o and a cocircuit b of N such that ω is in the closure of both o and b.

If ω is in the closure of some N-circuit, then every N-circuit is a $\{\omega\}$-circuit, and every N-cocircuit is an $\emptyset$-cocircuit. Since by Lemma 3.6 no Ψ-circuit ever meets a $\Psi^{\complement}$-cocircuit just once, we may apply Lemma 2.5 to deduce that $N = M_{\{\omega\}}(\mathcal{T})$. In the case that ω is not in the closure of any N-circuit a similar argument yields that $N = M_\emptyset(\mathcal{T})$. This completes the proof. $\square$

Having considered the case that the tree T is a ray, we now reduce the general case to this special case. Let N be a matroid with a tree decomposition (T, R) of adhesion 2. Let $Q = q_1, q_2, \ldots$ be a ray in T. We define R^Q to be the following coarsening of R. We define $R^Q_{q_i}$ to be the union of all the R_v such that in T the vertices v and q_i can be joined by a path that does not contain any other q_j.

Then (Q, R^Q) is a tree decomposition of N of adhesion 2. An N-circuit o has the end ω of Q in its closure with respect to (T, R) if and only if o has ω in its closure with respect to (Q, R^Q). So by Lemma 4.5, we deduce that there cannot be a circuit and a cocircuit of N that have a common end in both of their closures (with respect to (T, R)).
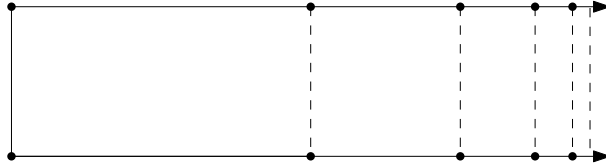
Figure 3: A non-nice tree of matroids

Let Ψ be the set of ends of T that appear in the closure of some circuit of N. Thus every N-circuit is a Ψ-circuit and every N-cocircuit is a Ψ^{C}-cocircuit. Since by Lemma 3.6 no Ψ-circuit ever meets a Ψ^{C}-cocircuit just once, we may apply Lemma 2.5 to deduce that $N = M_{\Psi}(\mathcal{T})$. Hence we get the following theorem.

Theorem 4.6. *Let N be a tame matroid with a tree decomposition (T, R) of adhesion 2.*

Then there is some $\Psi \subseteq \Omega(T)$ such that $N = M(\mathcal{T}(N, T, R), \Psi)$.

Combining this theorem with Theorem 3.10 yields:

Theorem 4.7. *Let N be a connected tame matroid. Then $N = M_{\Psi}(\mathcal{T})$ where each $M(t)$ is either a circuit, a cocircuit or else is 3-connected.*

Remark 4.8. *In the proof of this theorem, it might look as if we would have some freedom in choosing the set Ψ, namely that we could take Ψ to be any set containing all the ends to which some circuit converges and avoiding all ends to which some cocircuit converges. However, it can be shown that for every end in Ψ, there is a Ψ-circuit having this end in the closure.*

The arguments above make use of the fact that the matroid N is severly constrained by the restriction that each of its circuits comes from some precircuit of the tree, and each of its cocircuits comes from some precocircuit. In investigating how restrictive constraints of this form might be in general, we are led to the following question. Suppose that we have a tree $\mathcal{T} = (T, M)$ of matroids. We say a matroid N is a $\mathcal{T}$-matroid if every circuit of N is an $\Omega(T)$-circuit and every cocircuit of N is an $\emptyset$-cocircuit. How constrained is N? If $\mathcal{T}$ is not nice, then N can be quite unconstrained.

Example 4.9. Here the tree T is a ray and each $M(t) = M(C_4)$, arranged as in Figure 3. Then $M_{\emptyset}(\mathcal{T})$ is the free matroid but $M_{\Omega(T)}(\mathcal{T})$ consists of a single infinite circuit. So any pair of edges forms an $M_{\Omega(T)}(\mathcal{T})$-cocircuit which is not an $M_{\emptyset}(\mathcal{T})$-cocircuit.

However, if $\mathcal{T}$ is nice and N is tame then N has to be of the form $M_{\Psi}(\mathcal{T})$:

Theorem 4.10. *Let $\mathcal{T} = (T, M)$ be a nice tree of matroids of overlap 1, and let N be a tame $\mathcal{T}$-matroid. Then there is some $\Psi \subseteq \Omega(T)$ such that $N = M_{\Psi}(\mathcal{T})$.*
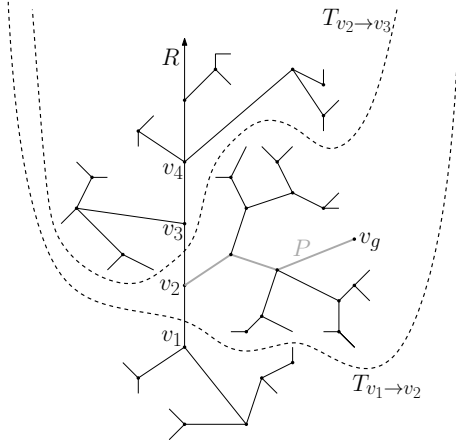
Figure 4: Objects appearing in the proof of Theorem 4.10

Proof. We begin by showing that there cannot be a circuit o and a cocircuit b of N such that there is some end ω of T in the closure of both o and b. So suppose for a contradiction that there are such o, b, and ω. We fix some notation, as illustrated in Figure 4. Pick a ray $R = v_1, v_2, \ldots$ in T to ω. By taking a suitable tail of R if necessary, we may assume that there is some edge f of b in $E(\mathcal{T}) \setminus E(\mathcal{T}_{v_1 \to v_2})$, some edge g of o in $E(\mathcal{T}_{v_1 \to v_2}) \setminus E(\mathcal{T}_{v_2 \to v_3})$ and some edge h of b in $E(\mathcal{T}_{v_2 \to v_3})$ (here we use that $\mathcal{T}$ is nice). Since $o \cap b$ is finite, we may even assume that no edge of $o \cap b$ lies in $E(\mathcal{T}_{v_1 \to v_2})$.

We may also assume that o has an edge in $E(\mathcal{T}_{v_2 \to v_1})$, so that both dummy edges of $E(v_2)$ are in $\hat{o}(v_2)$.

By Lemma 2.2 there is some $M_\emptyset(\mathcal{T})$-circuit o_b meeting b only in f and h. Let $(S, \hat{o})$ be an $\Omega(T)$-precircuit representing o, and $(S_b, \hat{o}_b)$ be an $\emptyset$-precircuit representing o_b. Let v_g be the node of T with $g \in E(v)$. Let P be the path joining v_2 to v_g in T. Let ∂ be the set of edges tt' of T with t in $V(P)$ but t' not in either $V(P)$ or $V(T_{v_2 \to v_3})$. For each edge $tt' \in \partial$ there is by niceness of $\mathcal{T}$ some $M_\emptyset(\mathcal{T}_{t \to t'})$-circuit $o_{t \to t'}$ through $e(tt')$. Let $(S_{t \to t'}, \hat{o}_{t \to t'})$ be an $M_\emptyset(\mathcal{T}_{t \to t'})$-precircuit representing $o_{t \to t'}$.

Now we build a $\emptyset$-precircuit $(S_C, \hat{C})$ from all this data as follows. First we set
$$S_C = (S_b \cap T_{v_2 \to v_3}) \cup P \cup \bigcup_{tt' \in \partial} S_{t \to t'}.$$

Then we take $\hat{C}(u)$ to be $\hat{o}_b(u)$ for $u \in (S_b \cap T_{v_2 \to v_3})$, $\hat{o}(u)$ for $u \in P$ and $\hat{o}_{t \to t'}(u)$ for $u \in S_{t \to t'}$. Let C be the underlying circuit of $(S_C, \hat{C})$. By Lemma 4.4, C is an N-circuit.

We now apply circuit elimination in N to the circuits o and C, eliminating the edge g and keeping the edge h. Call the resulting circuit C', and let $(S_{C'}, \hat{C}')$ be an $\Omega(T)$-precircuit representing C'. Let the vertices of P be, in order, $v_g =$

$p_1, p_2, \ldots p_k = v_2$. We shall show by induction on i that $p_i \notin S_{C'}$. For the base case, we note that if v_g were in $S_{C'}$ we would have to have $\hat{C}'(v_g) \subseteq \hat{o}(v_g) \setminus \{g\}$, which is impossible. For the induction step, we similarly note that if p_{i+1} were in $S_{C'}$ we would have to have $\hat{C}'(p_{i+1}) \subseteq \hat{o}(p_{i+1}) \setminus \{e(p_i p_{i+1})\}$, by the induction hypothesis, which is impossible. In particular, we deduce that $v_2 \notin S_{C'}$. On the other hand, we know that $h \in C'$ so that $S_{C'} \subseteq T_{v_2 \to v_3}$, so that $C' \cap b = \{h\}$, a contradiction.

We have now established that there cannot be a circuit o and a cocircuit b of N such that there is some end ω of T in the closure of both o and b. Let Ψ be the set of ends of T that appear in the closure of some circuit of N. Thus every N-circuit is a Ψ-circuit and every N-cocircuit is a $\Psi^{\complement}$-cocircuit. Since by Lemma 3.6 no Ψ-circuit ever meets a $\Psi^{\complement}$-cocircuit just once, we may apply Lemma 2.5 to deduce that $N = M_\Psi(\mathcal{T})$ as required. $\qquad\square$

5   Tame G-matroids are Ψ-matroids

Let G be a locally finite graph. Recall that a matroid N on the ground set $E(G)$ is a G-matroid if $\mathcal{C}(N) \subseteq \mathcal{C}(M_C(G))$ and $\mathcal{C}(N^*) \subseteq \mathcal{C}(M_{FC}(G)^*)$. Since $\mathcal{C}(M_{FC}(G)) \subseteq \mathcal{C}(M_C(G))$ and $\mathcal{C}(M_C^*(G)) \subseteq \mathcal{C}(M_{FC}^*(G))$ both $M_{FC}(G)$ and $M_C(G)$ are G-matroids, and an argument like that for Lemma 4.4 shows that for any G-matroid N we have $\mathcal{C}(M_{FC}(G)) \subseteq \mathcal{C}(N)$ and $\mathcal{C}(M_C^*(G)) \subseteq \mathcal{C}(N^*)$. The aim of this section is to prove the following.

Theorem 5.1. *Let G be a locally finite graph, and let N be a tame G-matroid. There there is some $\Psi \subseteq \Omega(G)$ such that $N = M_\Psi(G)$.*

For the rest of this section we fix some locally finite graph G and some tame G-matroid N.

In this section, we will have to use two different notions of path. Finite paths in graphs will simply be called *paths*, whereas paths in the topological sense, namely continuous images of the closed unit interval, will be called *topological paths*.

For a pair of points on a topological circle, there are two arcs joining them through the circle. To allow us to distinguish them, we shall make use of orientations of circles and topological paths. For distinct points x and y on an oriented circle $\vec{o}$ we use $x\vec{o}y$ to denote the (oriented) topological path from x to y through o whose orientation agrees with that of $\vec{o}$. We denote the other topological path by $x\overleftarrow{o}y$. If $x = y$, we do not take the trivial topological path but the topological path that goes all the way around the circle.

Lemma 5.2. *Let o be a topological circle in G. Let $v, w \in V(o)$ and let S be a finite set of vertices avoiding $V(o)$. Then for any orientation $\vec{o}$ of o, there is a finite v-w-path $P_{v,w}$ not meeting $v\vec{o}w$ in interior points and avoiding S.*

Moreover if $v\vec{o}w$ has at least two edges, then there is a bond b' of G that has $P_{v,w} \cup v\vec{o}w$ on one side and all interior vertices of $v\vec{o}w$ on the other side.

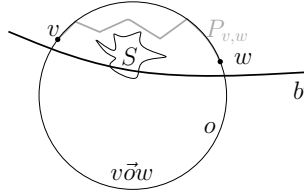Figure 5 gives an overview of the terminology used in this lemma.
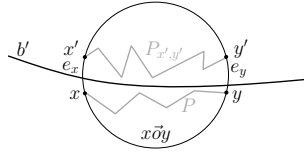
Figure 5: The situation of Lemma 5.2.

Figure 6: The situation of Lemma 5.3.

Proof. The proof is trivial if $v = w$. Thus we may assume that $v \neq w$. Let e_v be the first edge on $v\vec{o}w$ and e_w be the last edge on $v\vec{o}w$. Note that e_v and e_w exist since v and w are vertices. If $e_v = e_w$, we pick $P_{v,w} = e_v$. So we may assume that $e_v \neq e_w$.

Let $G' = G \setminus S$. Since o is a topological circle in $|G'|$, there is a finite bond b of G' meeting o in precisely e_v and e_w.

All edges and vertices of $v\vec{o}w - e_v - e_w - v - w$ are on the same side of b. Let C be the other side. Note that v and w are in C. Now let $P_{v,w}$ be some path in C joining v and w.

The bond b extends to a finite bond b' of G by adding finitely many deleted edges. The bond b' has the desired property. □

Lemma 5.3. *Let o be a circuit of the G-matroid N, and $\vec{o}$ be some orientation of o. Further, let $x, y \in V(o)$ and let $\vec{P}$ be an x-y-path meeting o in precisely x and y.*

Then $x\vec{o}y\overleftarrow{P}x \in \mathcal{C}(N)$.

Proof. The proof is trivial if $x = y$. Thus we may assume that $x \neq y$. Let e_x be the first edge on $x\vec{o}y$ and e_y be the last edge on $x\vec{o}y$. Let x' be the endvertex of e_x that is not x, and y' be the endvertex of e_y that is not y, as depicted in Figure 6.

Applying Lemma 5.2 to $\vec{o}$ with $S = V(P) - x - y$, yields an x'-y'-path $P_{x',y'}$, and a bond b' as in that lemma. By assumption the finite circuit $x\vec{P}ye_yy'\overleftarrow{P}_{x',y'}x'e_xx$ is an N-circuit. See Figure 6 to get an overview of all the definitions.

Now we apply circuit elimination in N to this new circuit and o eliminating e_x and keeping some $z \in E(P)$. Note that z exists since $x \neq y$. We obtain an

N-circuit $o' \subseteq (o - e_x) \cup P \cup P_{x',y'}$ including z.

It remains to show that $o' = x\vec{o}y \cup P$. Since each vertex of G is incident with 0 or 2 edges of o', we conclude that each edge adjacent to z on P is in o'. In fact an inductive argument yields that $P \subseteq o'$.

If $x\vec{o}y$ consists of a singe edge xy, then by the same argument xy also cannot be in o'. Thus $o' \subseteq x\vec{o}y\overleftarrow{P}x$. Since the latter is a topological cycle, we must have equality, hence $x\vec{o}y\overleftarrow{P}x \in \mathcal{C}(N)$. Thus we may assume that $x\vec{o}y$ includes at least two edges.

We know that o' is the union of P and some topological arc A from x to y. The edge set of this arc is included in $L := (o - e_x) \cup P_{x',y'}$. The set L meets b' precisely in e_y. Let K be the side of b' not containing x and y. Suppose for a contradiction that L includes an edge e_k from K. Then there are two disjoint arcs L_x and L_y from e_k to x and from e_k to y. By the Jumping-Arc Lemma [13, Lemma 8.5.3], both of these have to meet b', contradicting the fact that L meets b' just in e_y.

This means that $L \subseteq x\vec{o}y$. Since $x\vec{o}y$ is an x-y-arc, we actually get $L = x\vec{o}y$. Thus we have shown that $o' = x\vec{o}y \cup P$, which completes the proof. $\square$

Corollary 5.4. *Let $o \in \mathcal{C}(N)$, and $\vec{o}$ be some orientation of o. Further, let $x, y \in V(o)$ and let $\vec{P}$ be an x-y-path meeting $x\vec{o}y$ not in interior points.*
Then $x\vec{o}y\overleftarrow{P}x \in \mathcal{C}(N)$.

Remark 5.5. *The only difference between Lemma 5.3 and Corollary 5.4 is that in the second the path P may meet o in some of the interior points.*

Proof. We prove this by strong induction on $|P \cap o|$. Let z be the second point in the order of P in $P \cap o$ (the first such point is x). Now we apply Lemma 5.3 to o and xPz and obtain a new circuit $o_z := x\vec{o}z\overleftarrow{P}x$ and a new path $P_z := zPy$. Since $|o_z \cap P_z| < |o \cap P|$, we may apply the induction hypothesis. $\square$

Lemma 5.6. *Let $o \in \mathcal{C}(N)$ and let ω be an end of o. Then there is $o' \in \mathcal{C}(N)$ that has only the end ω in its closure.*

Proof. First we pick an orientation $\vec{o}$ of o. Then we pick a $\mathbb{Z}$-indexed family of distinct edges e_i such that their ordering on $\omega\vec{o}\omega$ is the same as the ordering of their indices and $(e_i|i > 0)$ and $(e_i|i < 0)$ both converge to ω.

Let s_i and t_i be the endvertices of e_i such that $s_i < t_i$ on $\omega\vec{o}\omega$. We repeatedly apply Corollary 5.4 to get a $\mathbb{Z}$-indexed family of vertex-disjoint t_i-s_{i+1}-paths P_i with P_i disjoint from $t_{i+1}\vec{o}s_i$.

Let D be the double ray obtained from sticking the P_i and the e_i together, formally:

$$\ldots t_{-1}P_{-1}s_0e_0t_0P_0s_1e_1t_1P_1\ldots$$

By construction, both tails of D belong to ω. So D is a topological cycle. It remains to show that $D \in \mathcal{C}(N)$. Suppose not, for a contradiction: Then $D \in \mathcal{I}(N)$, so there is a N-bond b with $b \cap D = \{e_0\}$.

Since N is tame, $o \cap b$ is finite, so there are only finitely many $i \in \mathbb{Z}$ such that b meets $t_i \vec{o} s_{i+1}$. Let K be the set of such i.

By applying Corollary 5.4 finitely often, we get a circuit o'' that meets b precisely in e_0. Formally,

$$o'' = o \setminus \left(\bigcup_{i \in K} (t_i \vec{o} s_{i+1}) \right) \cup \left(\bigcup_{i \in K} P_i \right)$$

So there are a circuit and a cocircuit of N which meet just once, which is the desired contradiction. $\square$

Lemma 5.7. *Let $b \in \mathcal{C}(N^*)$ and ω be an end in the closure of b. Assume there is a double ray $o \in \mathcal{C}(N)$ both of whose tails converge to ω.*

Then there is such an o that does not meet b.

Proof. We prove this by induction on $|o \cap b|$; the base case $|o \cap b| = 0$ is clear. The case $|o \cap b| = 1$ is impossible. So suppose for the induction step that $|o \cap b| \geq 2$. Thus we may pick $e, f \in o \cap b$. Since $b \in \mathcal{C}(M_{FC}(G))$, there is a finite circuit o' meeting b in precisely e and f. Note that o' is an N-circuit. Now pick z in the infinite component of $o \setminus o'$ containing ω.

Applying circuit elimination to o and o' eliminating e and keeping z yields an N-circuit $o'' \subseteq o \cup o' - e$ through z. By the choice of z, the subgraph with edge set $o \cup o' - e - z$ has two components one of which is a ray R from one endvertex of z converging to ω. Since each vertex is incident with 0 or 2 edges of o'' and $z \in o''$, the ray R must be included in o''. Hence o'' must be infinite, it also has only the end ω in its closure since $o'' \subseteq o \cup o' - e$ has no other end in its closure. Now $o'' \cap b \subseteq (o \cap b) - e$. This completes the induction step. $\square$

Lemma 5.8. *Let $b \in \mathcal{C}(N^*)$ and ω be an end in the closure of b.*

Then there is no double ray $o \in \mathcal{C}(N)$ both of whose tails converge to ω with $o \cap b = \emptyset$.

Proof. Suppose for a contradiction that there is such an N-circuit o. Let C_1 and C_2 be the two sides of b. Since $o \cap b = \emptyset$ and since the double ray o is connected as a subgraph, it lies entirely on one side, say C_1. Since G is locally finite, C_2 includes a ray R converging to ω.

Now we construct o-R-paths P_i as in Figure 7. Since o and R both have ω in their closure, there are infinitely many vertex disjoint R-o-paths $(P_i | i \in \mathbb{N})$. We enumerate the P_i such that in the linear order on R, the starting vertex of P_i is less than the starting vertex of P_j if and only if $i > j$. By Ramsey's theorem there is a tail R_o of o and $N \subseteq \mathbb{N}$ such that all P_i with $i \in N$ have their endvertex on R_o, and in the linear order on R_o, the endvertex of P_i is less than the endvertex of P_j if and only if $i > j$. By relabeling the indices of the P_i if necessary, we may assume that $N = \mathbb{N}$. Let s_i be the starting vertex of P_i, and t_i be its endvertex.

Now we prepare to apply the infinite circuit elimination axiom. We pick some edge x_i between t_{2i-1} and t_{2i} on R_o, and pick some $z \in o \setminus R_o$. Then
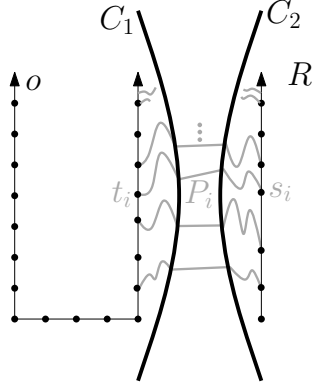
Figure 7: The o-R-paths P_i.

$C_{x_i} = t_{2i-1}R_o t_{2i}P_{2i}s_i R s_{2i-1}P_{2i-1}t_{2i-1}$ is a finite circuit. So C_{x_i} is an N-circuit. We apply circuit elimination to o and the C_{x_i} eliminating the x_i and keeping z. Thus there is an N-circuit o' through z that is included in:

$$\left(o \cup \left(\bigcup_{i \in \mathbb{N}} t_{2i}P_{2i}s_i R s_{2i-1}P_{2i-1}t_{2i-1} \right) \right) \setminus \{x_i | i \in \mathbb{N}\}$$

Since each vertex has degree 0 or 2 on o', no edge from any of the finite paths $X_i := t_{2i-1}R_o t_{2i}$ is in o'. Hence o' is included in:

$$D := (o \setminus R_o) \cup \left(\bigcup_{i \in \mathbb{N}} C_{x_i} \setminus X_i \right)$$

But D is a double ray. So $D = o'$ and is an N-circuit. But $D \cap b$ is infinite. This contradicts the tameness of N. □

Proof of Theorem 5.1. First we show that there cannot be an N-circuit o, and an N-cocircuit b that have a common end in their closure. Suppose for a contradiction there are such o and b. By Lemma 5.6, we get that there is such an o with only the end ω in its closure. By Lemma 5.7, we get there is such an o that additionally does not meet b. By Lemma 5.8, we then get the desired contradiction. So no end ω is ever in the closure of both a N-circuit and a N-cocircuit.

This motivates the following definition. Let Ψ be the set of ends that are in the closure of some N-circuit. Then every N-circuit is a Ψ-circuit and every N-cocircuit is a Ψ^{C}-cocircuit. Since N is a matroid, and the intersection of any Ψ-circuit with any Ψ^{C}-cocircuit is never of size 1 by Lemma 2.6, we are in a position to apply Lemma 2.5. Hence $N = M(G, \Psi)$. □

6 Outlook

Theorem 5.1 and Theorem 4.10 are closely related. Consider first Theorem 4.10, applied to a tree $\mathcal{T} = (T, M)$ of finite matroids. Then the matroid $M_\emptyset(\mathcal{T})$ is finitary and the matroid $M_{\Omega(T)}(\mathcal{T})$ is cofinitary. The niceness condition tells us even more: $M_\emptyset(\mathcal{T})$ is the finitarisation of $M_{\Omega(T)}(\mathcal{T})$, that is, its circuits are precisely the finite circuits of $M_{\Omega(T)}(\mathcal{T})$. Dually, $M_{\Omega(T)}(\mathcal{T})$ is the cofinitarisation of $M_\emptyset(\mathcal{T})$, that is, its cocircuits are precisely the finite cocircuits of $M_\emptyset(\mathcal{T})$. Theorem 4.10 determines the lattice of tame matroids N lying between these two in the sense that all their circuits are circuits of $M_{\Omega(T)}(\mathcal{T})$ and all their cocircuits are cocircuits of $M_\emptyset(\mathcal{T})$.

The situation in Theorem 5.1 is similar. Since G is a locally finite graph, $M_{FC}(G)$ is the finitarisation of $M_C(G)$ and $M_C(G)$ is the cofinitarisation of $M_{FC}(G)$. Once more, we have characterised the lattice of tame matroids lying between these two. The similarity runs even deeper. In [4] we show that Ψ-matroids for a locally finite graph G are naturally thought of as being constructed from trees of finite matroids associated to G (though these trees of matroids need not have overlap 1).

This leads us to consider the following general context. We denote the finitarisation of a matroid M by M^{fin} and the cofinitarisation of M by M^{cofin}. We say that matroids M_f and M_c are *twinned* if $M_f = M_c^{\mathrm{fin}}$ and $M_c = M_f^{\mathrm{cofin}}$. Then we say that N *lies between* M_f and M_c if all of its circuits are M_c-circuits and all of its cocircuits are M_f-cocircuits.

Open Question 6.1. *What is the structure of the lattice of tame matroids between M_f and M_c?*

Examples of twinned pairs (M_f, M_c) of matroids abound.

Proposition 6.2. *Let M be any finitary matroid. Then $((M^{\mathrm{cofin}})^{\mathrm{fin}})^{\mathrm{cofin}} = M^{\mathrm{cofin}}$, so that $(M^{\mathrm{cofin}})^{\mathrm{fin}}$ and M^{cofin} are twinned.*

Proof. No circuit of $(M^{\mathrm{cofin}})^{\mathrm{fin}}$ ever meets a cocircuit of M^{cofin} in just one element, so every cocircuit b of M^{cofin} is a coscrawl of $(M^{\mathrm{cofin}})^{\mathrm{fin}}$ by the dual of Lemma 2.4. The cocircuits in the union must all be cocircuits of $((M^{\mathrm{cofin}})^{\mathrm{fin}})^{\mathrm{cofin}}$ since b is finite. Similarly, every circuit of M is a scrawl of $(M^{\mathrm{cofin}})^{\mathrm{fin}}$, so that no cocircuit of $((M^{\mathrm{cofin}})^{\mathrm{fin}})^{\mathrm{cofin}}$ ever meets a circuit of M in just one element, and so every cocircuit b of $((M^{\mathrm{cofin}})^{\mathrm{fin}})^{\mathrm{cofin}}$ is a coscrawl of M^{cofin} by the dual of Lemma 2.4 applied to M and finiteness of b. Thus $((M^{\mathrm{cofin}})^{\mathrm{fin}})^{\mathrm{cofin}}$ and M^{cofin} have the same cocircuits. $\square$

Thus for example if we begin with the finite cycle matroid of a graph G then its cofinitarisation is twinned with the finitarisation of its cofinitarisation. In fact these are just the topological cycle matroid and finite cycle matroid of the finitely separable quotient of G, in which we identify any two vertices that cannot be separated by removing finitely many edges of G.

Another example arises when we glue together finite matroids along a tree $\mathcal{T}$ of arbitrary overlap. In such cases we need the matroids to be representable

over a common field in order to define the gluing: see [4] for details. Once more, if Ψ is a Borel set of ends then we get a matroid $M_\Psi(\mathcal{T})$ in which we only allow circuits to use ends from Ψ. Once more, $M_\emptyset(\mathcal{T})$ is finitary and $M_\Omega(\mathcal{T})$ is cofinitary. If the tree is nice then $M_\emptyset(\mathcal{T})$ and $M_\Omega(\mathcal{T})$ are twinned. We conjecture that in such cases all tame matroids between $M_\emptyset(\mathcal{T})$ and $M_\Omega(\mathcal{T})$ are of the form $M_\Psi(\mathcal{T})$.

Even though there are lots of examples, the lattices of tame matroids between twinned pairs of matroids are poorly understood. For example, if (M_f, M_c) is a twinned pair of matroids it is clear that every finite minor of M_f on a subset F of their common ground set also arises as a minor on the same subset F of any matroid lying between them. It isn't clear whether the converse holds, even if N has to be tame. Even the following question remains open:

Open Question 6.3. *We say a tame matroid is* binary *if it does not have $U_{2,4}$ as a minor [3]. Let (M_f, M_c) be a twinned pair of binary matroids. Must every tame matroid lying between them be binary?*

References

[1] H. Afzali and N. Bowler. Thin sums matroids and duality. Preprint 2012, current version available at
`http://www.math.uni-hamburg.de/spag/dm/papers/Thin_sums_and_duality.pdf`.

[2] E. Aigner-Horev, R. Diestel, and L. Postle. The structure of 2-separations of infinite matroids. Preprint 2012, current version available at
`http://www.math.uni-hamburg.de/home/diestel/papers/MatroidDecompositions.pdf`.

[3] N. Bowler and J. Carmesin. An excluded minors method for infinite matroids. Preprint 2012, current version available at
`http://www.math.uni-hamburg.de/spag/dm/papers/excludedminors_v5.pdf`.

[4] N. Bowler and J. Carmesin. Infinite matroids and determinacy of games. Preprint 2013, current version available at `arXiv:1301.5980 [math.CO]`.

[5] N. Bowler and J. Carmesin. Matroids with an infinite circuit-cocircuit intersection. Preprint 2012, current version available at
`http://www.math.uni-hamburg.de/spag/dm/papers/WildMatroids.pdf`.

[6] N. Bowler, J. Carmesin, and R. Christian. Infinite graphic matroids. In preparation.

[7] N. Bowler, J. Carmesin, and L. Postle. Reconstruction problems and wild matroids. In preparation.

[8] H. Bruhn and R. Diestel. Infinite matroids in graphs. arXiv:1011.4749 [math.CO], 2010.

[9] H. Bruhn, R. Diestel, M. Kriesell, R. Pendavingh, and P. Wollan. Axioms for infinite matroids. Advances in Mathematics, to appear; arXiv:1003.3919 [math.CO].

[10] Henning Bruhn and Maya Stein. Duality of ends. *cpc*, 19:47–60, 2010.

[11] William H. Cunningham and Jack Edmonds. A combinatorial decomposition theory. *Canad. J. Math.*, 32(3):734–765, 1980.

[12] R. Diestel. Locally finite graphs with ends: a topological approach. *Hamburger Beitr. Math.*, 340, 2009. see
`http://www.math.uni-hamburg.de/math/research/Preprints/hbm.html`.

[13] R. Diestel. *Graph Theory* (4th edition). Springer-Verlag, 2010.
Electronic edition available at:
`http://diestel-graph-theory.com/index.html`.

[14] R. Diestel and J. Pott. Dual trees must share their ends. Preprint 2011.

[15] A. Dress. Duality theory for finite and infinite matroids with coefficients. *Advances in Mathematics*, 59:97–123, 1986.

[16] George J. Minty. On the axiomatic foundations of the theories of directed linear graphs, electrical networks and network-programming. *J. Math. Mech.*, 15:485–520, 1966.

[17] J. Nikiel. Dendrons, dendritic spaces, and uniquely arcwise connected spaces. In *Papers on general topology and applications (Amsterdam, 1994)*, volume 788 of *Ann. New York Acad. Sci.*, pages 170–174. New York Acad. Sci., New York, 1996.

[18] J. Oxley. *Matroid Theory*. Oxford University Press, 1992.

[19] P. D. Seymour. Decomposition of regular matroids. *J. Combin. Theory Ser. B*, 28(3):305–359, 1980.

[20] C. Thomassen and A. Vella. Graph-like continua, augmenting arcs, and Menger's theorem. *Combinatorica*, 29, 2008. DOI: 10.1007/s00493-008-2342-9.