TREE AUTOMATA AND PIGEONHOLE CLASSES OF MATROIDS – II

DARYL FUNK, DILLON MAYHEW, AND MIKE NEWMAN

Abstract. Let ψ be a sentence in the monadic second-order logic of matroids. Let $\mathbb{F}$ be a finite field, and let $\mathcal{M}$ be the class of $\mathbb{F}$-representable matroids. Hliněný's Theorem says that there is a fixed-parameter tractable algorithm for testing whether matroids in $\mathcal{M}$ satisfy ψ, with respect to the parameter of branch-width. The main result of this paper extends Hliněný's Theorem by showing that such an algorithm also exists when $\mathcal{M}$ is the class of fundamental transversal matroids, lattice path matroids, bicircular matroids, or H-gain-graphic matroids, when H is a finite group.

1. Introduction

In the first paper of the series [6], we proved an extension of Hliněný's Theorem [8]. That theorem is concerned with testing sentences in the monadic second-order logic of matroids. This is the same logical language as used in [15]. Let ψ be a sentence in monadic second-order logic. Hliněný's Theorem says that there is a fixed-parameter tractable algorithm for testing whether matroids satisfy ψ, as long as the input class consists of matroids representable over a finite field $\mathbb{F}$. The input to a *fixed-parameter tractable* algorithm typically includes a numerical parameter, λ, and the running time is bounded by $O(f(\lambda)n^c)$, where n is the size of the input, c is a constant, and $f(\lambda)$ is a value that depends only on λ. In the case of Hliněný's Theorem, the parameter is the branch-width of the input matroid. Thus the theorem provides us with a polynomial-time algorithm when the input class is restricted to $\mathbb{F}$-representable matroids of bounded branch-width.

The main result of [6] is as follows.

Theorem 1.1. *Let $\mathcal{M}$ be a computably pigeonhole class of matroids. Let ψ be a sentence in MS_0. There is a fixed-parameter tractable algorithm for testing whether matroids in $\mathcal{M}$ satisfy ψ, where the parameter is branch-width.*

This sequel paper exploits Theorem 1.1 and related ideas to show that there is a fixed-parameter tractable algorithm for testing monadic sentences in other natural classes of matroids, beyond finite-field representable matroids. In particular, we show that Hliněný's Theorem can be extended as follows.

Date: October 14, 2019.

Theorem 1.2. *Let $\mathcal{M}$ be any of the following:*

(i) *the class of fundamental transversal matroids,*
(ii) *the class of lattice path matroids,*
(iii) *the class of bicircular matroids, or*
(iv) *the class of H-gain-graphic matroids, where H is a finite group.*

Let ψ be a sentence in MS_0. There is a fixed-parameter tractable algorithm for testing whether matroids in $\mathcal{M}$ satisfy ψ, where the parameter is branch-width.

We now explain computably pigeonhole matroid classes, along with some other associated concepts. Formal definitions are reserved for Section 3. Imagine that M is a matroid, and that U is a subset of $E(M)$. Let X and X' be subsets of U. Assume that $X \cup Z$ is independent if and only if $X' \cup Z$ is independent, for any subset $Z \subseteq E(M) - U$. We think of this as indicating that no subset of $E(M) - U$ can distinguish between X and X'. In this case we write $X \sim_U X'$. We put the elements of $E(M)$ into correspondence with the leaves of a ternary tree. If there are at most q equivalence classes under $\sim_U$ whenever U is a set displayed by an edge of the tree, then the *decomposition-width* of M is at most q. This notion of decomposition-width is equivalent to that used by Král [13] and by Strozecki [18, 19].

A class of matroids with bounded decomposition-width must have bounded branch-width [6, Corollary 2.8]. The converse does not hold (Lemma 4.1). Let $\mathcal{M}$ be a class of matroids, and assume that every subclass of $\mathcal{M}$ with bounded branch-width also has bounded decomposition-width. Then we say that $\mathcal{M}$ is a *pigeonhole* class of matroids. This is the case if and only if the dual class is pigeonhole ([6, Corollary 5.3]). The class of $\mathbb{F}$-representable matroids forms a pigeonhole class if and only if $\mathbb{F}$ is finite (Theorem 5.1 and Proposition 5.2). Fundamental transversal matroids (Theorem 6.3) and lattice path matroids also form pigeonhole classes (Theorem 7.2).

A stronger property arises quite naturally. Imagine that $\mathcal{M}$ is a class of matroids, that M is an arbitrary matroid in $\mathcal{M}$, and that U is an arbitrary subset of $E(M)$. Assume that whenever $\lambda_M(U)$, the connectivity value of U, is at most λ, there are at most $\pi(\lambda)$ equivalence classes under $\sim_U$, where $\pi(\lambda)$ is a value depending only on λ. In this case we say that $\mathcal{M}$ is *strongly pigeonhole* (Definition 3.4), and this implies that $\mathcal{M}$ is pigeonhole [6, Proposition 2.11]. The class of fundamental transversal matroids is strongly pigeonhole, and so is the class of $\mathbb{F}$-representable matroids when $\mathbb{F}$ is finite (Theorem 5.1). We do not know if any of the other classes in Theorem 1.2 are strongly pigeonhole, but we certainly believe this to be the case (Conjectures 9.1 and 9.3). In fact, we make the broad conjecture that the class of matroids that are transversal and cotransversal is a strongly pigeonhole class (Conjecture 9.2).

Theorem 1.1 relies on tree automata to test the sentence ψ, as does Hliněný's Theorem. These machines are described in Section 2. In order to construct a parse tree for the machine to process, we require a further strengthening of the pigeonhole property. It is not enough that there is a bound on the number of classes under $\sim_U$: we must be able to compute this equivalence relation efficiently. In fact, we are happy to compute a refinement of $\sim_U$, as long as this refinement does not have too many classes. If we are able to do this, then we say that the class is *computably pigeonhole* (Definition 3.8). Any computably pigeonhole class is also strongly pigeonhole. Matroids representable over a finite field (Theorem 5.1) are computably pigeonhole, and this gives us a proof of Hliněný's Theorem. The class of fundamental transversal matroids is also computably pigeonhole (Theorem 6.3).

In [6, Theorem 6.11] we proved that Theorem 1.1 holds under the weaker condition that the 3-connected members of $\mathcal{M}$ form a computably pigeonhole class. (However, we require that we can efficiently construct descriptions of minors, so the two theorems are independent of each other.) This was motivated by the fact that we do not know if bicircular matroids or H-gain-graphic matroids (H finite) form computably pigeonhole classes. (We conjecture this is the case in Conjecture 9.3). We have been able to show that the 3-connected bicircular matroids and the 3-connected H-gain-graphic matroids form computably pigeonhole classes (Theorem 8.4). This is then enough to establish cases (iii) and (iv) in Theorem 1.2.

Knowing that we have efficient model-checking for bicircular matroids gives us a new, and quite simple, proof of Courcelle's Theorem (Remark 8.6), which states that there is a fixed-parameter tractable algorithm for testing monadic second-order sentences in graphs, relative to the parameter of tree-width.

As well as proving positive results, we establish some negative propositions. Any class of matroids that contains the rank-3 uniform matroids and is closed under principal extensions is not pigeonhole (Corollary 4.2). Thus matroids representable over an infinite field are a non-pigeonhole class (Proposition 5.2). The class of transversal matroids is not pigeonhole, (Proposition 6.1) and nor is the class of gammoids (Remark 6.2). A different argument shows that the class of H-gain-graphic matroids is not pigeonhole when H is infinite (Proposition 8.8).

Oxley provides our reference for the basic concepts of matroid theory [16]. If M is a matroid, and (U, V) is a partition of $E(M)$, then $\lambda_M(U)$ is $r_M(U) + r_M(V) - r(M)$, and we call this the *connectivity value* of U. A *k-separation* is a partition, (U, V), of the ground set such that $|U|, |V| \geq k$, and $\lambda_M(U) < k$. A matroid is *n-connected* if it has no k-separations with $k < n$.

2. Tree automata

Definition 2.1. Let T be a tree with a distinguished *root* vertex, t. Assume that every vertex of T other than t has degree one or three, and that if T has more than one vertex, then t has degree two. The *leaves* of T are the degree-one vertices. In the case that t is the only vertex, we also consider t to be a leaf. Let $L(T)$ be the set of leaves of T. If T has more than one vertex, and v is a non-leaf, then v is adjacent with two vertices that are not contained in the path from v to t. These two vertices are the *children* of v. We distinguish the *left* child and the *right* child of v. Now let Σ be a finite *alphabet* of *characters*. Let σ be a function from $V(T)$ to Σ. Under these circumstances we say that (T, σ) is a Σ-*tree*.

Definition 2.2. A *tree automaton* is a tuple $(\Sigma, Q, F, \delta_0, \delta_2)$, where Σ is a finite alphabet, and Q is a finite set of *states*. The set of *accepting states* is a subset $F \subseteq Q$. The functions, $\delta_0 \colon \Sigma \to 2^Q$ and $\delta_2 \colon \Sigma \times Q \times Q \to 2^Q$, are *transition rules*.

Let $A = (\Sigma, Q, F, \delta_0, \delta_2)$ be an automaton and let (T, σ) be a Σ-tree with root t. We let $r \colon V(T) \to 2^Q$ be the function recursively defined as follows:

(i) if v is a leaf of T, then $r(v) = \delta_0(\sigma(v))$,
(ii) if v has left child v_L and right child v_R, then

$$r(v) = \bigcup_{(q_L, q_R) \in r(v_L) \times r(v_R)} \delta_2(\sigma(v), q_L, q_R).$$

We say that r is the *run* of the automaton A on (T, σ). Note that we define a union taken over an empty collection to be the empty set. We say that A *accepts* (T, σ) if $r(t)$ contains an accepting state.

Let i be a positive integer. Then $\{0,1\}^{\{i\}}$ denotes the set of functions from $\{i\}$ into $\{0,1\}$. Let Σ be a finite alphabet, and let (T, σ) be a Σ-tree. Let φ be a bijection from the finite set E into $L(T)$, and let Y_i be a subset of E. We construct the $(\Sigma \cup \Sigma \times \{0,1\}^{\{i\}})$-tree $\mathrm{enc}(T, \sigma, \varphi, \{Y_i\})$. The characters applied to the leaves of this tree will encode the subset Y_i. If v is a non-leaf vertex of T, then it receives the label $\sigma(v)$ in $\mathrm{enc}(T, \sigma, \varphi, \{Y_i\})$. However, if v is a leaf, then it receives a label $(\sigma(v), s)$, where $s \in \{0,1\}^{\{i\}}$ takes i to 1 if and only if $\varphi^{-1}(v)$ is in Y_i.

Definition 2.3. Let Σ be a finite set, and let A be a tree automaton with $\Sigma \cup \Sigma \times \{0,1\}^{\{i\}}$ as its alphabet. Let (T, σ) be a Σ-tree, and let φ be a bijection from the finite set E into $L(T)$. We define the set-system $M(A, T, \sigma, \varphi)$ as follows:

$$M(A, T, \sigma, \varphi) = (E, \{Y_i \subseteq E \colon A \text{ accepts } \mathrm{enc}(T, \sigma, \varphi, \{Y_i\})\}).$$

Now let Σ be a finite set, and let A be a tree automaton with alphabet $\Sigma \cup \Sigma \times \{0,1\}^{\{i\}}$. Let $M = (E, \mathcal{I})$ be a set-system. Assume there is a Σ-tree (T_M, σ_M), and a bijection $\varphi_M \colon E \to L(T_M)$ having the property that $M = M(A, T_M, \sigma_M, \varphi_M)$. In this case we say that (T_M, σ_M) is a *parse tree* for M (relative to the automaton A).

Note that if (T_M, σ_M) is a parse tree for M, then we can simulate an independence oracle for M by running A. We simply label each leaf v with the function taking i to 1 if $\varphi_M^{-1}(v)$ is in Y_i, and the function taking i to 0 if it is not. By then running A on the resulting tree, and testing to see if it accepts, we can check whether or not Y_i is in $\mathcal{I}$. This is idea is central to the proofs of Hliněný's Theorem and of Theorem 1.1.

3. Pigeonhole classes

This section states our central definitions: decomposition-width, pigeonhole classes, strongly pigeonhole classes, and computably pigeonhole classes. A *set-system* is a pair $(E,\mathcal{I})$ where E is a finite set and $\mathcal{I}$ is a family of subsets of E. We sometimes call the members of $\mathcal{I}$ the *independent* sets of the set-system.

Definition 3.1. Let $(E,\mathcal{I})$ be a set-system, and let U be a subset of E. Let X and X' be subsets of U. We say X and X' are *equivalent* (relative to U), written $X \sim_U X'$, if for every subset $Z \subseteq E - U$, the set $X \cup Z$ is in $\mathcal{I}$ if and only if $X' \cup Z$ is in $\mathcal{I}$.

Clearly $\sim_U$ is an equivalence relation on the subsets of U. No member of $\mathcal{I}$ is equivalent to a subset not in $\mathcal{I}$. When $\mathcal{I}$ is the set of independent sets of a matroid (more generally, when $\mathcal{I}$ is closed under subset containment), all dependent subsets of U are equivalent.

A *ternary tree* is one in which every vertex has degree three or one. A degree-one vertex is a *leaf*. Let $M = (E,\mathcal{I})$ be a set-system. A *decomposition* of M is a pair (T,φ), where T is a ternary tree, and φ is a bijection from E into the set of leaves of T. Let e be an edge joining vertices u and v in T. Then e partitions E into sets (U_e,V_e) in the following way: an element $x \in E$ belongs to U_e if and only if the path in T from $\varphi(x)$ to u does not contain v. We say that the partition (U_e,V_e) and the sets U_e and V_e are *displayed* by the edge e. Define $\mathrm{dw}(M;T,\varphi)$ to be the maximum number of equivalence classes in $\sim_U$, where the maximum is taken over all subsets, U, displayed by an edge in T. Define $\mathrm{dw}(M)$ to be the minimum value of $\mathrm{dw}(M;T,\varphi)$, where the minimum is taken over all decompositions (T,φ) of M. This value is then said to be the *decomposition-width* of M. If M is a matroid, then $\mathrm{dw}(M)$ is defined to be $\mathrm{dw}(E(M),\mathcal{I})$. Král [13] and Strozecki [18, 19] used an equivalent notion of decomposition-width.

Let M be a matroid. If (T,φ) is a decomposition of $M = (E(M),\mathcal{I}(M))$, then $\mathrm{bw}(M;T,\varphi)$ is the maximum value of

$$\lambda_M(U_e) + 1 = r_M(U_e) + r_M(V_e) - r(M) + 1,$$

where the maximum is taken over all partitions (U_e,V_e) displayed by edges of T. Now the *branch-width* of M (written $\mathrm{bw}(M)$) is the minimum value of $\mathrm{bw}(M;T,\varphi)$, where the minimum is taken over all decompositions of M. In [6, Corollary 2.8] we show that a class of matroids with bounded decomposition-width also has bounded branch-width. The converse is not true (see Lemma 4.1). This motivates the following definition.

Definition 3.2. Let $\mathcal{M}$ be a class of matroids. Then $\mathcal{M}$ is *pigeonhole* if, for every positive integer, λ, there is an integer $\rho(\lambda)$ such that $\mathrm{bw}(M) \le \lambda$ implies $\mathrm{dw}(M) \le \rho(\lambda)$, for every $M \in \mathcal{M}$.

So a class of matroids is pigeonhole if every subclass with bounded branch-width also has bounded decomposition-width. The next result is [6, Corollary 5.3].

Proposition 3.3. *Let $\mathcal{M}$ be a class of matroids. Then $\mathcal{M}$ is pigeonhole if and only if $\{M^*\colon M \in \mathcal{M}\}$ is pigeonhole.*

We often find that natural classes of matroids with the pigeonhole property also possess a stronger property.

Definition 3.4. Let $\mathcal{M}$ be a class of matroids. Assume that for every positive integer λ, there is a positive integer $\pi(\lambda)$, such that whenever $M \in \mathcal{M}$ and $U \subseteq E(M)$ satisfies $\lambda_M(U) \le \lambda$, there are at most $\pi(\lambda)$ equivalence classes under $\sim_U$. In this case we say that $\mathcal{M}$ is *strongly pigeonhole.*

In [6, Proposition 2.11], we give the easy proof that any class with the strong pigeonhole property also has the pigeonhole property.

Proposition 3.5. *The class of uniform matroids is strongly pigeonhole.*

Proof. Let M be a rank-r uniform matroid, and let U be a subset of $E(M)$ such that $\lambda_M(U) \le \lambda$, for some positive integer λ. Declare subsets $X, X' \subseteq U$ to be equivalent if:

(i) $|X|, |X'| > r_M(U)$,
(ii) $r_M(U) - \lambda < |X| = |X'| \le r_M(U)$, or
(iii) $|X|, |X'| \le r_M(U) - \lambda$.

Thus there are at most $\lambda + 2$ equivalence classes, and we will be done if we can show that this equivalence relation refines $\sim_U$. If $|X|, |X'| > r_M(U)$ then both X and X' are dependent, and hence they are equivalent under $\sim_U$. Since M is uniform, any subsets of U with the same cardinality will be equivalent under $\sim_U$. Therefore we need only consider the case that $|X|, |X'| \le r_M(U) - \lambda$. Assume that $Z \subseteq E(M) - U$, and $X \cup Z$ is independent while $X' \cup Z$ is dependent. Since $X' \cup Z$ is dependent, it follows that $|X' \cup Z| > r(M)$. As $X \cup Z$ is independent, we see that $|Z| \le r_M(E(M) - U)$. Therefore

$$r(M) < |X' \cup Z| = |X'| + |Z| \le r_M(U) - \lambda + r_M(E(M) - U).$$

Hence $r_M(U) + r_M(E(M) - U) - r(M) > \lambda$, and we have a contradiction to $\lambda_M(U) \le \lambda$. $\square$

Theorem 1.1 is concerned with matroid algorithms. For the purposes of measuring the efficiency of these algorithms, we restrict our attention to matroid classes where there is a succinct representation, such as graphic matroids or finite-field representable matroids.

Definition 3.6. Let $\mathcal{M}$ be a class of matroids. A *succinct representation* of $\mathcal{M}$ is a relation, Δ, from $\mathcal{M}$ into the set of finite binary strings. We write $\Delta(M)$ to indicate any string in the image of $M \in \mathcal{M}$. We insist that there is a polynomial p and a Turing Machine which, when given any input $(\Delta(M), X)$, where $M \in \mathcal{M}$ and X is a subset of $E(M)$, will return an answer to the question "Is X independent in M?" in time bounded by $p(|E(M)|)$.

Note that the length of the string $\Delta(M)$ is no more than $p(|E(M)|)$. A graph provides a succinct representation of a graphic matroid, and a matrix provides a succinct representation of a finite-field representable matroid.

Definition 3.7. Let Δ be a succinct representation of $\mathcal{M}$, a class of matroids. We say that Δ is *minor-compatible* if there is a polynomial-time algorithm which will accept any tuple $(\Delta(M), X, Y)$ when $M \in \mathcal{M}$ and X and Y are disjoint subsets of $E(M)$, and return a string of the form $\Delta(M/X\backslash Y)$.

The proof of Theorem 1.1 proceeds by constructing a tree automaton which tests whether an MS_0 sentence is satisfied by the input matroid. In order to construct parse trees for the automaton to process, we need to be able to efficiently compute the equivalence classes of $\sim_U$. In fact, we are happy to compute an equivalence relation that refines $\sim_U$, as long as it does not have too many classes.

Definition 3.8. Let $\mathcal{M}$ be a class of matroids with a succinct representation Δ. Assume there is a constant, c, and that for every integer, $\lambda > 0$, there is an integer, $\pi(\lambda)$, and a Turing Machine, M_λ, with the following properties: M_λ takes as input any tuple of the form $(\Delta(M), U, X, X')$, where M is in $\mathcal{M}$, $U \subseteq E(M)$ satisfies $\lambda_M(U) \leq \lambda$, and X and X' are subsets of U. The machine M_λ computes an equivalence relation, $\approx_U$, on the subsets of U, so that M_λ accepts $(\Delta(M), U, X, X')$ if and only if $X \approx_U X'$. Furthermore,

(i) $X \approx_U X'$ implies $X \sim_U X'$,
(ii) the number of equivalence classes under $\approx_U$ is at most $\pi(\lambda)$, and
(iii) M_λ runs in time bounded by $O(\pi(\lambda)|E(M)|^c)$.

Under these circumstances, we say that $\mathcal{M}$ is *computably pigeonhole* (relative to Δ).

Clearly a computably pigeonhole class of matroids is also strongly pigeonhole.

4. Non-pigeonhole classes

Next we develop some tools for proving negative results. We want to certify that certain classes are not pigeonhole. Let G be a simple graph with vertex set edge set $\{e_1, \ldots, e_m\}$ and $\{v_1, \ldots, v_n\}$ where $n \geq 3$. We define $m(G)$ to be the rank-3 sparse paving matroid with ground set $\{v_1, \ldots, v_n\} \cup \{e_1, \ldots, e_m\}$. The only non-spanning circuits of $m(G)$ are the sets $\{v_i, e_k, v_j\}$, where e_k is an edge of G joining the vertices v_i and v_j.

Lemma 4.1. *Let $\mathcal{M}$ be a class of matroids. Assume there are arbitrarily large integers, N, such that $\mathcal{M}$ contains a matroid isomorphic to $m(K_N)$. Then $\mathcal{M}$ contains rank-3 matroids with arbitrarily high decomposition-width. Hence $\mathcal{M}$ is not pigeonhole.*

Proof. Observe that rank-3 matroids have branch-width at most four, so if $\{M \in \mathcal{M} \colon r(M) = 3\}$ has unbounded decomposition-width, then $\mathcal{M}$ is certainly not pigeonhole. Assume for a contradiction that every rank-3 matroid in $\mathcal{M}$ has decomposition-width at most K.

Let n be a positive integer. Erdős and Rado [5] proved that there is a least integer $\phi(n, k)$, such that in any collection of distinct n-element sets with at least $\phi(n, k)$ members, there is a subcollection of k sets having a single pairwise intersection. Thus if a simple graph has at least $\phi(2, k)$ edges, it has either a vertex of degree at least k, or a matching containing at least k edges. Abbott, Hanson, and Sauer [1] proved that $\phi(2, k)$ is $k(k-1)$ when k is odd, and $(k-1)^2 + (k-2)/2$ when k is even. Thus we can choose $k > K$ such that $k^2 \geq \phi(2, k)$.

Next we choose an integer N such that

$$\frac{1}{3}\left(N + \binom{N}{2}\right) \geq 7k^2 + 2k$$

and $\mathcal{M}$ contains a matroid, M, isomorphic to $m(K_N)$. By relabelling, we assume that the ground set of M is $\{v_1, \ldots, v_N\} \cup \{e_{ij} \colon 1 \leq i < j \leq N\}$ and the only non-spanning circuits are of the form $\{v_i, e_{ij}, v_j\}$. Let (T, φ) be a decomposition of M with the property that if U is any displayed set, then $\sim_U$ has at most K classes. Using [16, Lemma 14.2.2], we choose an edge e in T such that each of the displayed sets, U_e and V_e, contains at least

$$\frac{1}{3}|E(M)| = \frac{1}{3}\left(N + \binom{N}{2}\right)$$

elements. Let G be a complete graph with vertex set $\{v_1, \ldots, v_N\}$ and edge set $\{e_{ij} \colon 1 \leq i < j \leq N\}$, where e_{ij} joins v_i to v_j. We colour a vertex or edge red if it belongs to U_e, and blue if it belongs to V_e.

Assume that there are at least $2k$ red vertices and at least $2k$ blue vertices. Without loss of generality we can assume that there is a matching in G consisting of k red edges, each of which joins a red vertex to a blue vertex. Thus we can find elements $v_{i_1}, \ldots, v_{i_k}$ in U_e and elements $v_{j_1}, \ldots, v_{j_k}$ in V_e such that $e_{i_p j_p}$ is in U_e for each p. If p and q are distinct, then $\{v_{i_p}, e_{i_p j_p}, v_{j_p}\}$ is a circuit of M while $\{v_{i_q}, e_{i_q j_q}, v_{j_p}\}$ is a basis. Hence $\{v_{i_p}, e_{i_p j_p}\}$ and $\{v_{i_q}, e_{i_q j_q}\}$ are inequivalent under $\sim_{U_e}$. This means that $\sim_{U_e}$ has at least k equivalence classes. As $k > K$, this is a contradiction, so we assume without loss of generality that there are fewer than $2k$ red vertices.

Assume some red vertex is joined to at least k blue vertices by red edges. Then there is an element $v_i \in U_e$ and elements $v_{j_1}, \ldots, v_{j_k} \in V_e$ such that e_{ij_p} is in U_e for each p. For distinct p and q, we see that $\{v_i, e_{ij_p}, v_{j_p}\}$ is a circuit while $\{v_i, e_{ij_q}, v_{j_p}\}$ is a basis. Therefore $\{v_i, e_{ij_p}\}$ and $\{v_i, e_{ij_q}\}$ are inequivalent under $\sim_{U_e}$. We again reach the contradiction that there are at least k equivalence classes under $\sim_{U_e}$. Now we can deduce that there are fewer than $2k^2$ red edges that join a red vertex to a blue vertex.

There are fewer than $2k$ red vertices and fewer than

$$\binom{2k}{2} < 4k^2$$

red edges that join two red vertices. Since the number of red edges and vertices is at least one third of $N + \binom{N}{2}$, we see that the number of red edges joining two blue vertices is at least

$$\frac{1}{3}\left(N + \binom{N}{2}\right) - (2k + 2k^2 + 4k^2) \geq k^2 \geq \phi(2, k).$$

Therefore the subgraph induced by such red edges contains either a vertex of degree at least k, or a matching containing at least k edges.

In the former case, there are elements $v_i, v_{j_1}, \ldots, v_{j_k} \in V_e$ such that e_{ij_p} is in U_e for each p. Then $\{v_i, e_{ij_p}, v_{j_p}\}$ is a circuit, and $\{v_i, e_{ij_p}, v_{j_q}\}$ is a basis for distinct p and q, so $\{v_i, v_{j_p}\}$ and $\{v_i, v_{j_q}\}$ are inequivalent under $\sim_{V_e}$. This leads to a contradiction, so there is a matching of at least k edges. Therefore we can find elements $v_{i_1}, \ldots, v_{i_k}, v_{j_1}, \ldots, v_{j_k}$ in V_e such that each $e_{i_pj_p}$ is in U_e. For distinct p and q, we see that $\{v_{i_p}, e_{i_pj_p}, v_{j_p}\}$ is a circuit and $\{v_{i_q}, e_{i_pj_p}, v_{j_q}\}$ is a basis. Therefore $\{v_{i_p}, v_{j_p}\}$ and $\{v_{i_q}, v_{j_q}\}$ are inequivalent under $\sim_{V_e}$, so we reach a final contradiction that completes the proof. □

Let F be a flat of the matroid M. Let M' be a single-element extension of M, and let e be the element in $E(M') - E(M)$. We say that M' is a *principal extension* of M by F if $F \cup e$ is a flat of M' and whenever $X \subseteq E(M)$ spans e in M', it spans $F \cup e$.

Corollary 4.2. *Let $\mathcal{M}$ be a class of matroids. If $\mathcal{M}$ contains all rank-3 uniform matroids, and is closed under principal extensions, then it is not pigeonhole.*

Proof. We note that $m(K_N)$ can be constructed by starting with a rank-3 uniform matroid, the elements of which represent the vertices of K_N. The elements representing edges are then added via principal extensions. The result now follows from Lemma 4.1. □

5. Representable matroids

The next result is not surprising, and has been utilised by both Hliněný [8] and Král [13].

Theorem 5.1. *Let $\mathbb{F}$ be a finite field. The class of $\mathbb{F}$-representable matroids is computably pigeonhole.*

Proof. Assume that $|\mathbb{F}| = q$. Let $\mathcal{M}$ be the class of $\mathbb{F}$-representable matroids. We consider the succinct representation Δ that sends each matroid in $\mathcal{M}$ to an $\mathbb{F}$-matrix representing it. Let M be a rank-r matroid in $\mathcal{M}$, and let U

be a subset of M. We use V to denote $E(M) - U$. We identify M with a multiset of points in the projective geometry $P = \mathrm{PG}(r-1, q)$ (we lose no generality in assuming that M is loopless). If X is a subset of $E(M)$, then $\langle X \rangle$ will denote its closure in P.

Assume that $\lambda_M(U) \leq \lambda$. Grassman's identity tells us that the rank of $\langle U \rangle \cap \langle V \rangle$ is equal to $r(U) + r(V) - r(M) \leq \lambda$. We define the equivalence relation $\approx_U$ so that if X and X' are subsets of U, then $X \approx_U X'$ if both X and X' are dependent, or both are independent and $\langle X \rangle \cap \langle V \rangle = \langle X' \rangle \cap \langle V \rangle$. Deciding whether $X \approx_U X'$ holds requires only elementary linear algebra, and it can certainly be accomplished in time bounded by $O(|E(M)|^c)$ for some constant c. Since $\langle U \rangle \cap \langle V \rangle$ is a subspace of P with affine dimension at most $\lambda - 1$, it contains at most $(q^\lambda - 1)/(q-1)$ points. Therefore $2^{q^{\lambda-1}+q^{\lambda-2}+\cdots+1} + 1$ is a crude upper bound on the number of $(\approx_U)$-classes. It remains only to show that $\approx_U$ refines $\sim_U$.

Assume that $X \approx_U X'$, and yet $X \cup Z$ is independent while $X' \cup Z$ is dependent, where Z is a subset of V. Then X is independent, so X' is independent also. Let C be a circuit contained in $X' \cup Z$. As both X' and Z are independent, neither $X' \cap C$ nor $Z \cap C$ is empty. Now the rank of $\langle X' \cap C \rangle \cap \langle Z \cap C \rangle$ is

$$r(X' \cap C) + r(Z \cap C) - r(C) = |X' \cap C| + |Z \cap C| - (|C| - 1) = 1.$$

Let c be the point of P that is in $\langle X' \cap C \rangle \cap \langle Z \cap C \rangle$. Since c is in $\langle X' \rangle \cap \langle V \rangle$, our assumption tells us it is also in $\langle X \rangle \cap \langle V \rangle$.

Assume c is not in X. Since it is in $\langle X \rangle$, we can let C_X be a circuit contained in $X \cup c$ that contains c. If c is in Z, then $X \cup Z$ contains C_X, and we have a contradiction, so c is not in Z. We let C_Z be a circuit contained in $Z \cup c$ that contains c. Circuit elimination between C_X and C_Z shows that $X \cup Z$ contains a circuit, and again we have a contradiction. Therefore c is in X. If c is not in Z, then $Z \cup c \subseteq X \cup Z$ contains a circuit. Therefore c is in Z. As X and Z are disjoint subsets of $E(M)$, but c is identified with elements of both, we conclude that M contains a parallel pair, with one element in X, and the other in Z. Again $X \cup Z$ is dependent, and we have a final contradiction. □

Hliněný's Theorem [8] follows immediately from Theorems 1.1 and 5.1. We note that proofs of Hliněný's Theorem can also be derived from the works by Král [13] and Strozecki [19].

Proposition 5.2. *Let $\mathbb{K}$ be an infinite field. Then the class of $\mathbb{K}$-representable matroids is not pigeonhole.*

Proof. This follows almost immediately from Corollary 4.2 and [14, Lemma 2.1]. □

6. Transversal matroids

Proposition 6.1. *The class of transversal matroids is not pigeonhole.*

Proof. By Proposition 3.3, we can prove that the class of transversal matroids is not pigeonhole by proving the same statement for the class of cotransversal matroids. Certainly this class contains all rank-3 uniform matroids. Recall that the matroid M is cotransversal if and only if it is a *strict gammoid* [9]. This means that there is a directed graph G with vertex set $E(M)$, and a distinguished set, T, of vertices, where $X \subseteq E(M)$ is independent in M if and only if there are $|X|$ vertex-disjoint directed paths, each of them starting with a vertex in X and terminating with a vertex in T. Assume that G is such a directed graph, and that F is a flat of M. Create the graph G' by adding the new vertex e, and arcs directed from e to each of the vertices in F. It is an easy exercise to verify that if M' is the strict gammoid corresponding to G', then M' is a principal extension of M by F. This demonstrates that the class of cotransversal matroids is closed under principal extensions, so the proposition follows by Corollary 4.2. □

Remark 6.2. From Proposition 6.1 we see that any class of matroids containing transversal matroids is not pigeonhole. In particular, the class of gammoids is not pigeonhole.

In contrast to Proposition 6.1, in subsequent sections we will show that three subclasses of transversal matroids are pigeonhole: fundamental transversal matroids (Theorem 6.3), lattice path matroids (Theorem 7.2), and bicircular matroids (Theorem 8.4).

6.1. **Fundamental transversal matroids.** Transversal matroids can be thought of geometrically as those obtained by placing points freely on the faces of a simplex. A transversal matroid is *fundamental* if there is a point placed on each vertex of that simplex. More formally, a transversal matroid is fundamental if it has a basis, B, such that $r(B \cap Z) = r(Z)$, for every cyclic flat Z (see [3]). From this it is easy to see that the dual of a fundamental transversal matroid is also fundamental.

Let G be a bipartite graph, with bipartition $A \cup B$. There is a fundamental transversal matroid, $M[G]$, with $A \cup B$ as its ground set, where $X \subseteq A \cup B$ is independent if and only if there is a matching, M, of G such that $|M| = |X \cap A|$ and each edge in M joins a vertex in $X \cap A$ to a vertex in $B - X$. In this case we say that M *certifies* X to be independent. This definition implies that B is a basis of $M[G]$, and $r(B \cap Z) = r(Z)$ for any cyclic flat Z. Moreover, any fundamental transversal matroid can be represented in this way. The transversal matroid on the ground set $A \cup B$ represented by this bipartite graph is equal to $M[G]$. Note that we can represent $M[G]$ with a standard bipartite presentation by adding an auxiliary vertex, b', for each vertex $b \in B$, and making b' adjacent only to b. We then swap the labels on b and b'.

Theorem 6.3. *The class of fundamental transversal matroids is computably pigeonhole.*

Proof. We consider the succinct representation of fundamental transversal matroids that involves representing such a matroid with a bipartite graph. Let $M[G]$ be a fundamental transversal matroid, where $A \cup B$ is a bipartition of the bipartite graph G, and B is a basis of $M[G]$. Let (U, V) be a partition of $A \cup B$, and assume that $\lambda_{M[G]}(U) \leq \lambda$. Let H be the subgraph of G induced by edges that join vertices in $B \cap U$ to vertices in $A \cap V$, and vertices in $B \cap V$ to vertices in $A \cap U$.

Claim 6.3.1. Any matching of H contains at most λ edges.

Proof. Let M be a matching in H. Let A_U and A_V, respectively, be the set of vertices in $A \cap U$ (respectively $A \cap V$) that are incident with an edge in M. Therefore $|A_U| + |A_V| = |M|$. If we restrict M to edges incident with vertices in $A \cap U$, then it certifies that $(B \cap U) \cup A_U$ is an independent subset of U. Similarly, $(B \cap V) \cup A_V$ is an independent subset of V. Therefore

$$\lambda \geq r(U) + r(V) - r(M[G]) \geq |B \cap U| + |A_U| + |B \cap V| + |A_V| - |B| = |M|. \quad \square$$

We can find a maximum matching of H, using one of a number of polynomial-time algorithms. It follows from Kőnig's Theorem [12] that H contains a vertex cover, S, such that $|S| \leq \lambda$. Furthermore, Kőnig's Theorem is constructive: given a maximum matching of H, we can find S in polynomial time. From this point onwards, we regard S as being fixed.

Let X be an independent subset of U, and let M be a matching that certifies its independence. We will construct a *signature*, $\mathcal{C}(X, M)$. Signatures of subsets of V will be defined symmetrically, so in fact we let $\{P, Q\}$ be $\{U, V\}$, and we let X be an independent subset of P, with M a matching certifying the independence of X. Recall that this means that $|M| = |X \cap A|$ and each edge of M joins a vertex in $X \cap A$ to a vertex in $B - X$. The signature $\mathcal{C}(X, M)$ is a sequence $(S_1, \mathcal{S}_2, S_3, S_4)$, where S_1, S_3, and S_4 are subsets of $B \cap P \cap S$, $A \cap P \cap S$, and $B \cap Q \cap S$, respectively, and $\mathcal{S}_2$ is a collection of subsets of $A \cap Q \cap S$. We define $\mathcal{C}(X, M)$ as follows.

(i) S_1 is the set of vertices in $B \cap P \cap S$ that are either in X or incident with an edge in M.
(ii) A subset $Z \subseteq A \cap Q \cap S$ is in $\mathcal{S}_2$ if and only if there is a matching M' satisfying $M \subseteq M'$ and $|M' - M| = |Z|$, where each edge in $M' - M$ joins a vertex in Z to a vertex in $(B \cap P) - (S \cup X)$. Note that $\mathcal{S}_2$ is closed under subset inclusion.
(iii) S_3 is the set of vertices in $A \cap P \cap S$ that are joined by an edge of M to a vertex in $(B \cap Q) - S$.
(iv) S_4 is the set of vertices in $B \cap Q \cap S$ that are joined by an edge in M to a vertex in $A \cap P$.

We illustrate these definitions in Figure 1. This shows a graph, G, with bipartition $A \cup B$, and a partition, (P, Q), of $A \cup B$. The edges not in H cross the diagram diagonally, and are drawn with dashed lines, while the unbroken edges are the edges of H. In this example the vertex cover, S, contains nine vertices, which are marked with squares. Observe that every edge of H is incident with a vertex in S. The set $X \subseteq P$ is marked by filled vertices. Its independence is certified by the matching M, which is drawn with heavy lines. Vertices in the sets S_1, S_3, and S_4 are marked. The family $\mathcal{S}_2$ contains the empty set, and the singleton set that contains the vertex marked S_2.

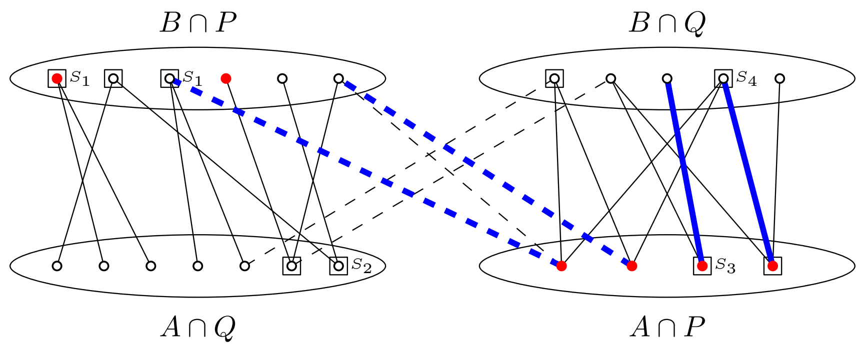

Figure 1. Defining a signature.

Claim 6.3.2. Let X be an independent subset of P. Let (S_1, Z, S_3, S_4) be a sequence of sets from $B \cap P \cap S$, $A \cap Q \cap S$, $A \cap P \cap S$, and $B \cap Q \cap S$. We can test in polynomial time whether there is a matching M, certifying the independence of X, such that $\mathcal{C}(X, M) = (S_1, \mathcal{S}_2, S_3, S_4)$ where Z is in $\mathcal{S}_2$.

Proof. To start with, we check that S_1 contains $X \cap B \cap P \cap S$ and that S_3 is contained in X. If this is not the case, then we halt and return the answer NO, so now we assume that $X \cap B \cap P \cap S \subseteq S_1$ and $S_3 \subseteq X$.

Our strategy involves constructing an auxiliary graph, G', by deleting vertices and edges from G. The construction of G' is best described by the diagram in Figure 2. Any vertex not shown in this diagram is deleted in the construction of G'. Thus from $B \cap P$ we delete any vertex in X, and any vertex in $(B \cap P \cap S) - S_1$. From $A \cap Q$ we delete any vertex not in Z. From $A \cap P$ we delete those vertices not in X. Note that the assumption in the first paragraph of this proof means that we have not deleted any vertex in S_3. In $B \cap Q$, we delete those vertices in $(B \cap Q \cap S) - S_4$.

Next we delete any edge of G that is not represented by an edge in Figure 2. For example, we delete any edge joining a vertex in S_3 to a vertex outside of $(B \cap Q) - S$. This completes the description of G', which can obviously be constructed in polynomial time.

6.3.2.1. The following statements are equivalent:

(i) There is a matching, M, of G, such that M certifies the independence of X and $\mathcal{C}(X, M) = (S_1, \mathcal{S}_2, S_3, S_4)$ with $Z \in \mathcal{S}_2$.

(ii) G' has a matching incident with every vertex in $(X \cup Z) \cap A$ and $(S_1 - X) \cup S_4$.

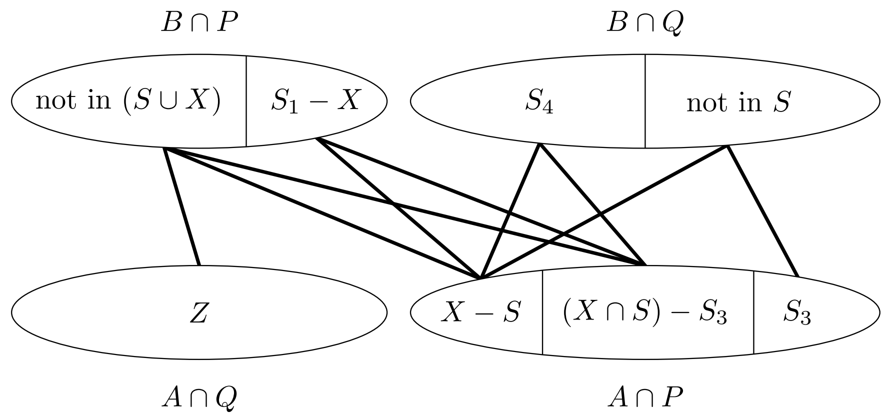

FIGURE 2. The construction of G'.

Proof. Assume (i) holds. Then $|M| = |X \cap A|$, and every vertex in $X \cap A$ is incident with an edge of M. Let M' be a matching such that $M \subseteq M'$, $|M' - M| = |Z|$, and each edge of $M' - M$ joins a vertex in Z to a vertex in $(B \cap P) - (S \cup X)$. Then every edge in $M' - M$ is an edge of G'. Let $ab \in M$ be an edge joining $a \in A$ to $b \in B$. Then a is in X. Note b is not in X, for no edge of M is incident with a vertex in $X \cap B$. If b is in $B \cap P \cap S$, then it is in S_1, by definition. So b is not in $(B \cap P \cap S) - S_1$. If b is in $B \cap Q \cap S$, then it is in S_4, so b is not in $(B \cap Q \cap S) - S_4$. Therefore b is a vertex of G'. This is enough to show that if a is in $X - S$, then ab is an edge of G'. Now assume a is in $(X \cap S \cap A) - S_3$. The previous discussion shows that the only way ab can fail to be an edge of G' is if b is in $(B \cap Q) - S$. But in this case, a would be in S_3, a contradiction. Finally, assume that a is in S_3. Then the definition of S_3 means that b is in $(B \cap Q) - S$, and again ab is an edge of G'. Thus we have shown that M' is a matching of G'. Every vertex in $(X \cup Z) \cap A$ is incident with an edge in M', and the same statement is true for vertices in $(S_1 - X) \cup S_4$, as $\mathcal{C}(X, M) = (S_1, \mathcal{S}_2, S_3, S_4)$. Therefore (ii) holds.

Now assume (ii) holds. Let M' be a matching of G' such that each vertex in $(X \cup Z) \cap A$ or $(S_1 - X) \cup S_4$ is incident with an edge of M'. Let M be the set of edges in M' incident with vertices in $X \cap A$. There is no vertex in $X \cap B$ contained in G', so it immediately follows that in G, M certifies the independence of X. Every vertex in $S_1 - X$ is incident with an edge of M, and no vertex of $(B \cap P \cap S) - S_1$ is (since these vertices are not in G). Therefore in $\mathcal{C}(X, M)$, the first entry is S_1, as desired. Every edge of $M' - M$ joins a vertex in Z to a vertex in $(B \cap P) - (S \cup X)$, so M' certifies that Z belongs to the second entry of $\mathcal{C}(X, M)$. Any vertex in S_3 is matched by M to a vertex in $(B \cap Q) - S$, and no vertex in $(X \cap S \cap A) - S_3$ is, by the construction of G'. Finally, every vertex in S_4 is matched by M to a vertex in $A \cap P$, and no vertex of $(B \cap Q \cap S) - S_4$ is (since these vertices are not in G'). Therefore $\mathcal{C}(X, M) = (S_1, \mathcal{S}_2, S_3, S_4)$, where Z is in $\mathcal{S}_2$, so (i) holds. $\square$

Now we complete the proof of Claim 6.3.2. To test whether M exists, we find a maximum-sized matching of G', using standard methods. If this matching is incident with all the vertices in $(X \cup Z) \cap A$ (and is thus complete), then we continue, otherwise we return NO. So now assume that the vertices in $(X \cup Z) \cap A$ are all matched. It is easy to see that we can use alternating-path methods to test whether there is a matching that matches all the vertices in $(S_1 - X) \cup S_4$ as well as those in $(X \cup Z) \cap A$. We return YES if such a complete matching exists, and NO otherwise, observing that 6.3.2.1 justifies the correctness of this algorithm. $\square$

For any independent subset $X \subseteq P$, let $\mathcal{C}(X)$ be the set

$$\{\mathcal{C}(X, M)\colon M \text{ is a matching certifying that } X \text{ is independent}\}.$$

Now we define the equivalence relation $\approx_U$. If X and X' are subsets of U, then say that $X \approx_U X'$ if both X and X' are dependent, or both are independent and $\mathcal{C}(X) = \mathcal{C}(X')$. Note that the number of certificates is at most the number of families of subsets of S, namely 2^{2^λ}. Therefore the number of $(\approx_U)$-classes is no more than $2^{2^{2^\lambda}} + 1$. To test whether $X \approx_U X'$, we first test whether X and X' are independent. We can certainly test this in polynomial-time via a standard matching algorithm. Assuming both X and X' are independent, we simply go through each possible certificate, and check that each certificate belongs to $\mathcal{C}(X)$ if and only if it belongs to $\mathcal{C}(X')$. According to Claim 6.3.2, we can accomplish this in time bounded by $O(2^{2^\lambda}|\Delta(M)|^c)$, for some constant c.

Now our final task in the proof of Theorem 6.3 is to show that $\approx_U$ refines $\sim_U$. To this end, we assume that $X \subseteq U$ and $Y \subseteq V$ are independent subsets of $M[G]$. Let $S_X = (S_1, \mathcal{S}_2, S_3, S_4)$ be a signature in $\mathcal{C}(X)$, and let $T_Y = (T_1, \mathcal{T}_2, T_3, T_4)$ be a member of $\mathcal{C}(Y)$. We declare S_X and T_Y to be *compatible* if the following conditions hold:

(i) $S_1 \cap T_4 = \emptyset$,
(ii) $T_3 \in \mathcal{S}_2$,
(iii) $S_3 \in \mathcal{T}_2$, and
(iv) $S_4 \cap T_1 = \emptyset$.

The remainder of the proof will follow immediately from Claim 6.3.3, and its converse (Claim 6.3.4).

Claim 6.3.3. Let $X \subseteq U$ and $Y \subseteq V$ be independent subsets of $M[G]$. If $X \cup Y$ is independent in $M[G]$ then there are signatures $S_X \in \mathcal{C}(X)$ and $T_Y \in \mathcal{C}(Y)$ such that S_X and T_Y are compatible.

Proof. Let M be a matching certifying that $X \cup Y$ is independent. Then no edge of M is incident with a vertex in $(X \cup Y) \cap B$. Let M_X and M_Y be the subsets of M consisting of edges incident with vertices in X (respectively Y). We assert that the signatures $\mathcal{C}(X, M_X)$ and $\mathcal{C}(Y, M_Y)$ are compatible.

Let $\mathcal{C}(X, M_X)$ be $(S_1, \mathcal{S}_2, S_3, S_4)$ and let $\mathcal{C}(Y, M_Y)$ be $(T_1, \mathcal{T}_2, T_3, T_4)$. Then S_1 is the set of vertices in $B \cap U \cap S$ that are either in X, or incident with an edge of M_X. On other hand, T_4 is the set of vertices in $B \cap U \cap S$ that are joined by an edge of M_Y to a vertex in $A \cap V$. No edge in M_Y is incident with an edge in M_X, or with a vertex in $B \cap X$, so it is clear that S_1 and T_4 are disjoint. Similarly, S_4 is the set of vertices in $B \cap V \cap S$ that are joined by an edge of M_X to a vertex in $A \cap U$, and T_1 is the set of vertices in $B \cap V \cap S$ that are either in Y, or incident with a vertex in M_Y. This implies that $S_4 \cap T_1 = \emptyset$.

Note that T_3 is the set of vertices in $A \cap V \cap S$ that are joined by an edge of M_Y to a vertex in $(B \cap U) - S$. Let M' be the union of M_X along with the set of edges in M_Y that are incident with a vertex in T_3. Clearly M' is a matching as it is a subset of M. Also, $M_X \subseteq M'$ and $|M' - M_X| = |T_3|$. Each edge in $M' - M_X$ is incident with a vertex in T_3, and with a vertex in $(B \cap U) - S$. Furthermore, no such edge is incident with a vertex in X, since edges of M join vertices in $(X \cup Y) \cap A$ to vertices in $B - (X \cup Y)$. Therefore each edge in $M' - M_X$ joins a vertex of T_3 to one in $(B \cap U) - (S \cup X)$. We have established that T_3 is contained in $\mathcal{S}_2$. A similar argument shows that S_3 is in $\mathcal{T}_2$. Therefore $\mathcal{C}(X, M_X)$ and $\mathcal{C}(Y, M_Y)$ are compatible, as we claimed. □

Claim 6.3.4. Let $X \subseteq U$ and $Y \subseteq V$ be independent subsets of $M[G]$. If there are signatures $S_X \in \mathcal{C}(X)$ and $T_Y \in \mathcal{C}(Y)$ such that S_X and T_Y are compatible, then $X \cup Y$ is independent in $M[G]$.

Proof. We assume that $\mathcal{C}(X, M_X) = (S_1, \mathcal{S}_2, S_3, S_4)$ and $\mathcal{C}(Y, M_Y) = (T_1, \mathcal{T}_2, T_3, T_4)$ are compatible signatures. We will construct a matching that certifies the independence of $X \cup Y$.

Recall that S_3 is the subset of $A \cap U \cap S$ containing vertices that are joined by edges of M_X to vertices in $(B \cap V) - S$. Let M''_X be the subset of M_X containing edges that are incident with vertices in S_3. Since S_3 is in $\mathcal{T}_2$, there is a matching, M'_Y, such that $M_Y \subseteq M'_Y$, $|M'_Y - M_Y| = |S_3|$, and each edge of $M'_Y - M_Y$ joins a vertex in S_3 to one in $(B \cap V) - (S \cup Y)$. Similarly, we let M''_Y be the subset of M_Y containing edges that are incident with vertices in T_3. Thus each edge in M''_Y joins a vertex in T_3 to a vertex in $(B \cap U) - S$. As T_3 is in $\mathcal{S}_2$, we can let M'_X be a matching such that $M_X \subseteq M'_X$, $|M'_X - M_X| = |T_3|$, and each edge of $M'_X - M_X$ joins a vertex in T_3 to a vertex in $(B \cap U) - (S \cup X)$. We now make the definition

$$M = (M'_X - M''_X) \cup (M'_Y - M''_Y).$$

We will prove that M is a matching certifying the independence of $X \cup Y$.

6.3.4.1. M is a matching.

Proof. If not then, there is a vertex w, and distinct edges $wx \in M'_X - M''_X$ and $wy \in M'_Y - M''_Y$. In the first case, we assume that w is in A. Assume also that w is in U. No edge of M_Y is incident with a vertex in $A \cap U$. Therefore wy is in $M'_Y - M_Y$. This means that wy joins a vertex of S_3 to a vertex in $(B \cap V) - (S \cup Y)$. In particular this means that w is in S_3. No edge in $M'_X - M_X$ is incident with a vertex in $A \cap U$, so wx is not in $M'_X - M_X$. Therefore it is in M_X, so wx is an edge of M_X that is incident with a vertex in S_3. But this means that wx is in M''_X, so we have a contradiction. If w is in V, then we reach the similar contradiction that wy is in M''_Y. Therefore we must now assume that w is in B.

We assume that w is in $B \cap U$. Any edge in M'_Y that is incident with a vertex in $B \cap U$ is also incident with a vertex in $A \cap V$. Therefore y belongs to $A \cap V$. If w is not in S, then y is in S, for otherwise wy is an edge of H that is not incident with the vertex cover S. But in this case, y is in T_3, so wy belongs to M''_Y, and we have a contradiction. Thus w is in S. If wx is in $M'_X - M_X$, then wx joins a vertex in T_3 to a vertex in $(B \cap U) - (S \cup X)$. This is impossible, as we have already confirmed that w is in S. Hence wx is in M_X, so w is in S and is incident with an edge of M_X, meaning that it is in S_1. Furthermore, the edge wy means that w is in T_4. Thus $S_1 \cap T_4 \neq \emptyset$, and we have a contradiction to the fact that $\mathcal{C}(X, M_X)$ and $\mathcal{C}(Y, M_Y)$ are compatible. If w is in V, then we reach the symmetric contradiction that either wx is in M''_X, or w is in $S_4 \cap T_1$. This completes the proof that M is a matching. $\square$

Let ab be an edge of $M'_X - M''_X$ with $a \in A$ and $b \in B$. We wish to demonstrate that b is not in $X \cup Y$. If ab is in $M'_X - M_X$, then ab joins a vertex in T_3 to a vertex in $(B \cap U) - (S \cup X)$. In this case, b is certainly not in X. Since b is in $B \cap U$, and $Y \subseteq V$, it follows that b is also not in Y. Therefore we will now assume that ab is not in $M'_X - M_X$, and thus ab is in M_X. Each edge of M_X joins a vertex of $A \cap X$ to a vertex of $B - X$, so b is not in X. Since $X \subseteq U$, it follows that a is in $A \cap U$. Assume that b is in Y, so that it belongs to $B \cap V$. If b is not in S, then a is in S, for otherwise ab is an edge of H that is not incident with the vertex cover S. In this case ab is an edge of M_X joining a vertex in $A \cap U \cap S$ to a vertex in $(B \cap V) - S$, so a is in S_3, and ab is in M''_X, a contradiction. Therefore b is in S. As b is in $Y \cap S$, it follows that it is in T_1. But the edge ab certifies that b is in S_4. Therefore $S_4 \cap T_1 \neq \emptyset$, and we have a contradiction to the fact that $\mathcal{C}(X, M_X)$ and $\mathcal{C}(Y, M_Y)$ are compatible. We have shown that b is not in $X \cup Y$, and a symmetrical argument shows that no edge of $M'_Y - M''_Y$ is incident with a vertex in $B \cap (X \cup Y)$. Hence every edge of M joins a vertex of $A \cap (X \cup Y)$ to a vertex in $B - (X \cup Y)$.

Let w be a vertex in $A \cap X$. If w is not incident with an edge of $M'_X - M''_X$, then it is incident with an edge of M''_X, and hence w is in S_3. But in this case w is incident with an edge of $M'_Y - M_Y$. By symmetry, we now see that every vertex in $A \cap (X \cup Y)$ is incident with an edge in M, so M certifies the independence of $X \cup Y$, exactly as we desired. $\square$

Let X and X' be two independent subsets of U such that $X \approx_U X'$. Then $\mathcal{C}(X) = \mathcal{C}(X')$. Let $Y \subseteq V$ be an independent set such that $X \cup Y$ is independent. Claim 6.3.3 shows that there are compatible signatures $S_X \in \mathcal{C}(X)$ and $T_Y \in \mathcal{C}(Y)$. As S_X is also in $\mathcal{C}(X')$ it follows from Claim 6.3.4 that $X' \cup Y$ is independent. This implies that $X \sim_U X'$, so $\approx_U$ refines $\sim_U$, as desired. Now the proof of Theorem 6.3 is complete. □

Case (i) in Theorem 1.2 follows immediately from Theorem 6.3 and Theorem 1.1.

Remark 6.4. Theorem 6.3 shows that although a class of matroids may be strongly pigeonhole, its minor-closure may not even be pigeonhole. We can deduce this from Remark 6.2 because the smallest minor-closed class containing the fundamental transversal matroids is the class of gammoids.

7. Lattice path matroids

The class of *lattice path* matroids was introduced by Bonin, de Mier, and Noy [2]. It is closed under duality and minors [2, Theorems 3.1 and 3.5]. Although we have not succeeded in proving the class to be computably pigeonhole, we do show that it is pigeonhole (Theorem 7.2), and we describe an algorithm that constructs a parse tree for a given lattice path matroid (Theorem 7.3). Combined with results from [6], this shows that there is a fixed-parameter tractable algorithm for testing MS_0 sentences in lattice path matroids.

Let m and r be integers, and let P and Q be strings composed of r copies of N and m copies of E. Any such string is identified with a path in the integer lattice from $(0,0)$ to (m,r) using North and East steps. We insist that P never goes above Q, so that for any initial substring in P, the number of N steps does not exceed the number of N steps in the corresponding substring of Q. The matroid $M[P,Q]$ has $\{1,\ldots,m+r\}$ as its ground set. An *intermediate path* is a string composed of r copies of N and m copies of E that does not go above Q or below P. Note that P and Q are both intermediate paths. Let $L = l_1 l_2 \cdots l_{m+r}$ be an intermediate path, where each l_i is either N or E, and let $N(L)$ be $\{i \colon l_i = N\}$. Then the family of bases of $M[P,Q]$ is $\{N(L) \colon L \text{ is an intermediate path}\}$.

Let $G[P,Q]$ be the graph whose vertices are those lattice points in $\mathbb{Z}^2$ that appear in an intermediate path. If (i,j) and (i',j') are two such lattice points, then they are adjacent in $G[P,Q]$ if and only if $|i-i'|+|j-j'| = 1$. Let e be an edge in $G[P,Q]$, and assume that e joins (i,j) to either $(i+1,j)$ or $(i,j+1)$. In this case we define $d(e)$ to be $i+j+1$. A *staircase* is a set $\{e \in E(G[P,Q]) \colon d(e) = k\}$, where $1 \leq k \leq m+r$. Figure 3 shows the staircase of edges with $d(e) = 7$. We identify the elements in $\{1,\ldots,m+r\}$ with the staircases of $G[P,Q]$.

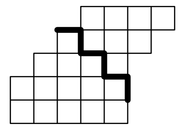

FIGURE 3. A staircase in a lattice path presentation.

Proposition 7.1. *Let $M = M[P, Q]$ be a lattice path matroid, and assume that* $\mathrm{bw}(M) \leq \lambda$. *No staircase of $G[P, Q]$ is incident with more than $3\lambda - 1$ vertices.*

Proof. Assume otherwise. Then there is an odd number o such that $o > 3\lambda - 2$ and there is a staircase with at least o vertices. So let q be an integer such that $2q + 1 > 3\lambda - 2$ and there is a staircase with at least $2q + 1$ vertices. It is straightforward to see that this implies the existence of a $q \times q$ square of the integer lattice contained in the region bounded by P and Q. This in turn implies that $M[P, Q]$ has a $U_{q,2q}$-minor, by [10, Lemma 4.1]. The branch-width of this minor is $\lceil 2q/3 \rceil + 1$ by [16, Exercise 14.2.5]. Since $\mathrm{bw}(M) \leq \lambda$, we deduce that $\lceil 2q/3 \rceil + 1 \leq \lambda$ by [16, Proposition 14.2.3]. From this it follows that

$$\begin{array}{rcl} \lceil 2q/3 \rceil & \leq & \lambda - 1 \\ \Rightarrow \quad 2q/3 & \leq & \lambda - 1 \\ \Rightarrow \quad 2q + 1 & \leq & 3\lambda - 2 \end{array}$$

and we have a contradiction. □

Theorem 7.2. *The class of lattice path matroids is pigeonhole.*

Proof. Let λ be a positive integer. We must show that the class of lattice path matroids with branch-width at most λ has bounded decomposition-width. Let $M = M[P, Q]$ be a lattice path matroid, where P and Q are paths from $(0, 0)$ to (m, r), and assume that $\mathrm{bw}(M) \leq \lambda$. We construct a decomposition of M by starting with path of $m + r - 2$ vertices, adjoining a leaf to each internal vertex, and two leaves to each end-vertex of the path. This describes the tree T. Let φ be the bijection from $\{1, \ldots, m + r\}$ that labels the leaves of T in a linear way, so that the only sets displayed by the decomposition are singleton sets, complements of singleton sets, and sets of the form $\{1, \ldots, i\}$ or $\{i + 1, \ldots, m + r\}$.

Let (U_e, V_e) be a partition displayed by an edge e in T. If $|U_e| = 1$ or $|V_e| = 1$, then it is clear that $\sim_{U_e}$ has at most two equivalence classes. Therefore we will assume that $U_e = \{1, \ldots, i\}$, and show that $\sim_{U_e}$ has at most $2^{3\lambda - 1} + 1$ equivalence classes. This will complete the proof. (The case where $U_e = \{i + 1, \ldots, m + r\}$ is essentially identical.)

Consider the graph $G[P, Q]$, and let R be the path consisting of the edges, x, satisfying $d(x) = i$. Let $w_1, \ldots, w_t$ be the vertices in R, starting from the top-left corner of the path. Proposition 7.1 implies that $t \leq 3\lambda - 1$. Let X be an independent subset of U_e. Let $\mathcal{C}(X)$ be the subset of $\{w_1, \ldots, w_t\}$

such that w_j is in $\mathcal{C}(X)$ if and only if there is an intermediate path, L, such that $X \subseteq N(L)$, and the last vertex of L in the path R is w_j. We declare $X, X' \subseteq U_e$ to be equivalent if they are both dependent, or if they are both independent and $\mathcal{C}(X) = \mathcal{C}(X')$. There are at most $2^{3\lambda-1}+1$ classes in this relation, so we will be done if we can show that this equivalence refines $\sim_{U_e}$. Let X and X' be equivalent subsets of U_e. If both are dependent then obviously $X \sim_{U_e} X'$, so assume that X and X' are independent and $\mathcal{C}(X) = \mathcal{C}(X')$. Let $X \cup Z$ be independent for some subset $Z \subseteq V_e$. Let L be an intermediate path such that $X \cup Z \subseteq N(L)$. Let w_j be the last vertex of L to appear in R. We can let L' be an intermediate path such that $X' \subseteq N(L')$, and the last vertex of L' to appear in R is w_j. If we concatenate the first segment of L' up to w_j, and the segment of L appearing after w_j, then we obtain an intermediate path that certifies the independence of $X' \cup Z$. This shows that $X \sim_{U_e} X'$, so the proof is complete. □

Theorem 7.3. *There is a fixed-parameter tractable algorithm (with respect to the parameter of branch-width) which takes as input any lattice path matroid and produces a parse tree for that matroid.*

Proof. We consider the succinct representation of lattice path matroids via the paths P and Q. Let $M = M[P,Q]$ be a lattice path matroid, where P and Q are paths from $(0,0)$ to (m,r), and assume that $\mathrm{bw}(M) = \lambda$. We construct the tree T_M as shown in Figure 4. The bijection φ_m takes the element $\ell \in \{1,\ldots,m+r\}$ to the leaf v_ℓ.

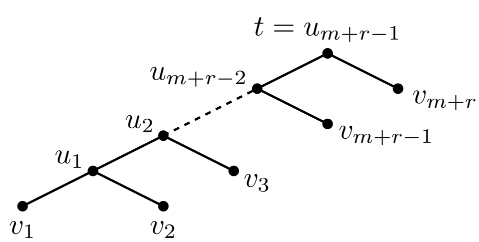

FIGURE 4. The parse tree for a lattice path matroid.

The labelling σ_M applies a function to each node of T_M. The function applied to any leaf is the identity function on $\{0,1\}$. Next consider the function, f, applied to u_ℓ. Let R and R', respectively, be the paths in $G[P,Q]$ consisting of edges, x, satisfying $d(x) = \ell$ in the first case, and $d(x) = \ell+1$ in the second. Let the vertices in R and R' be $w_1,\ldots,w_s$ and $w'_1,\ldots,w'_t$, starting in the top-left corners of these paths. Note that $s,t \leq 3\lambda-1$ by Proposition 7.1. The domain of f will be $\{1,\ldots,s,\mathsf{dep}\} \times \{0,1\}$. If the input includes dep, then the output is also dep. Otherwise, we consider the input $(j,1)$. The output of f is k, where w'_k is the vertex of R' reached from w_j by one north step, assuming this vertex exists. If it does not, then a north step from w_j takes us out of the region bounded by P and Q, and in this case we define $f(j,1) = \mathsf{dep}$. The output of f on $(j,0)$ is k, where w'_k is reached from w_j via an east step, assuming that w'_k exists. If it does not, then there must be a vertex w'_k one step north of w_j, and we define $f(j, 0)$ to be k. Now we have completed the description of the tree (T_M, σ_M).

Define the set Σ to contain the identity function on $\{0, 1\}$, and all functions from $\{1, \ldots, s, \mathsf{dep}\} \times \{0, 1\}$ into $\{1, \ldots, t, \mathsf{dep}\}$, where $s, t \leq 3\lambda - 1$. The automaton A has $\Sigma \cup \Sigma \times \{0, 1\}^{\{i\}}$ as its alphabet. The state space of A is $\{1, \ldots, 3\lambda - 1, \mathsf{dep}\}$, and the only non-accepting state is dep. The transition rule δ_0 operates on (f, s) by taking it to $\{f(s(i))\}$, whenever f is the identity function on $\{0, 1\}$ and s is in $\{0, 1\}^{\{i\}}$. Any other input is taken to dep by δ_0. If f is a function from $\{1, \ldots, s, \mathsf{dep}\} \times \{0, 1\}$ to $\{1, \ldots, t, \mathsf{dep}\}$, then $\delta_2(f, j, k)$ is $\{f(j, k)\}$ whenever (j, k) is in the domain of f. Any other output of δ_2 is dep. This completes the description of A.

Let Y_i be a subset of $\{1, \ldots, m + r\}$. The run of A on $\mathrm{enc}(T_M, \sigma_M, \varphi_M, \{Y_i\})$ is very easy to understand. As it processes the tree, A constructs a path starting at $(0, 0)$, considering each staircase in turn. If the current staircase is in Y_i, then A appends a north step to the path. If the staircase is not in Y_i, it appends an east step when it is able to do so without going outside the region bounded by P and Q, and otherwise appends a north step. If the constructed path goes outside this region, the automaton switches to the state dep and remains there. Otherwise, the state applied to a node records the last vertex in the constructed path by giving its position in the current staircase. Thus the states applied by the run of A either tell us that Y_i is dependent, or record an intermediate path, L, such that Y_i is contained in $N(L)$. Thus (T_M, σ_M) is a parse tree relative to A. □

The next result follows immediately from [6, Proposition 6.1] and Theorem 7.3. It establishes case (ii) in Theorem 1.2.

Theorem 7.4. *Let ψ be any sentence in MS_0. There is a fixed-parameter tractable algorithm for testing whether lattice path matroids satisfy ψ, where the parameter is branch-width.*

8. Frame matroids

Let G be a graph with edge set E. We allow G to contain loops and parallel edges. If X is a subset of E, we use $G[X]$ to denote the subgraph with edge set X, containing exactly those vertices that are incident with an edge in X. Similarly, if N is a set of vertices, then $G[N]$ is the induced subgraph of G with N as its vertex set. A *theta subgraph* consists of two distinct vertices joined by three internally-disjoint paths. A *linear class* of cycles in G is a family, $\mathcal{B}$, of cycles such that no theta subgraph of G contains exactly two cycles in $\mathcal{B}$. Let $\mathcal{B}$ be a linear class of cycles in G. A cycle in $\mathcal{B}$ is *balanced*, and a cycle not in $\mathcal{B}$ is *unbalanced*. A subgraph of G is *unbalanced* if it contains an unbalanced cycle, and is otherwise *balanced*.

Frame matroids were introduced by Zaslavsky [21]. The frame matroid, $M(G, \mathcal{B})$, has E as its ground set. The circuits of $M(G, \mathcal{B})$ are the edge sets of balanced cycles, and the edge sets of minimal connected subgraphs containing at least two unbalanced cycles, and no balanced cycles. Such a subgraph is either a theta subgraph or a *handcuff*. A *tight handcuff* contains two edge-disjoint cycles that have exactly one vertex in common. A *loose handcuff* consists of two vertex-disjoint cycles and a path having exactly one vertex in common with each of the two cycles. Note that if $\mathcal{B}$ contains every cycle, then $M(G, \mathcal{B})$ is a graphic matroid. The set $X \subseteq E$ is independent in $M(G, \mathcal{B})$ if and only if $G[X]$ contains no balanced cycle, and each connected component of $G[X]$ contains at most one cycle. The rank of X in $M(G, \mathcal{B})$ is the number of vertices in $G[X]$, minus the number of balanced components of $G[X]$.

Proposition 8.1. *Let $M = M(G, \mathcal{B})$ be a 3-connected frame matroid, and let (U, V) be a partition of the edge set of G such that $\lambda_M(U) \leq \lambda$. There are at most $14\lambda - 12$ vertices that are incident with edges in both U and V.*

Proof. Let n be the number of vertices in G. Let n_U and n_V be the number of vertices in $G[U]$ and $G[V]$, respectively. Let N be the set of vertices that are in both $G[U]$ and $G[V]$, so $n + |N| = n_U + n_V$. Each vertex in N is incident with a connected component of $G[U]$, and with a connected component of $G[V]$. Since G is connected, each component of $G[U]$ or $G[V]$ contain at least one vertex of N. Thus the connected components of $G[U]$ induce a partition of N. There are no coloops in M, and it follows that if a component of $G[U]$ contains only a single, non-loop, edge, then that edge joins two vertices of N. Let a be the number of such components. Next we claim that if X is a connected component of $G[U]$ such that X is balanced and contains at least two edges, then X contains at least three vertices of N. If this is not true, then we can easily verify that M has a 1- or 2-separation, contradicting the hypotheses of the theorem. Assume that there are b balanced components of $G[U]$ with more than one edge, and let $\alpha_i, \ldots, \alpha_b$ be the numbers of vertices these components share with N. Our claim shows that $\alpha_i \geq 3$ for each i. Finally, assume there are c unbalanced components in $G[U]$, and these components intersect N in $\beta_1, \ldots, \beta_c$ vertices, respectively. Thus $|N| = 2a + \sum \alpha_i + \sum \beta_i$, and $r_M(U) = n_U - (a + b)$.

Let x be the number of components of $G[V]$ consisting of a single non-loop edge. Assume there are y balanced components of $G[V]$ with more than one edge, and that these intersect N in $\gamma_1, \ldots, \gamma_y$ vertices. Let z be the number of unbalanced components of $G[V]$, and assume that they intersect N in $\delta_1, \ldots, \delta_z$ vertices, respectively. So we have $|\gamma_i| \geq 3$, $|N| = 2x + \sum \gamma_i + \sum \delta_i$, and $r_M(V) = n_V - (x + y)$. Because G is connected, $r(M) \geq n - 1$, and $r(M) = n - 1$ if and only if G is balanced. Now we observe that

$$\begin{aligned} \lambda \geq r_M(U) + r_M(V) - r(M) &\geq n_U + n_V - (a + b + x + y) - (n - 1) \\ &= |N| - (a + b + x + y) + 1. \end{aligned}$$

This last quantity is equal to $a + \sum \alpha_i + \sum \beta_i - (b + x + y) + 1$, and also to $x + \sum \gamma_i + \sum \delta_i - (a + b + y) + 1$, so both are at most λ. By adding the two inequalities together, we obtain

$$2\lambda \geq \sum \alpha_i + \sum \beta_i + \sum \gamma_i + \sum \delta_i - 2(b+y) + 2.$$

But because each α_i is at least three, we also have $b \leq \frac{1}{3}\sum \alpha_i$, and symmetrically $y \leq \frac{1}{3}\sum \gamma_i$. Therefore

$$6(\lambda - 1) \geq \sum \alpha_i + 3\sum \beta_i + \sum \gamma_i + 3\sum \delta_i. \tag{1}$$

The edges counted by a form a matching. Therefore they are an independent set in M. As $r_M(U) + r_M(V) - r(M) \leq \lambda$, submodularity tells us that the intersection of $\mathrm{cl}_M(U)$ and $\mathrm{cl}_M(V)$ has rank at most λ. Thus there are at least $a - \lambda$ components of $G[U]$ that consist of a single, non-loop, edge that is not in $\mathrm{cl}_M(V)$. No such edge can be incident with one of the components of $G[V]$ counted by x, for this would mean that a vertex of G has degree equal to two, implying that M contains a series pair. This is impossible, since M is 3-connected (and we can obviously assume that it has more than three elements). Nor can such an edge join two vertices counted by the variables $\delta_1, \ldots, \delta_z$, for then the edge joins two components of $G[V]$ that contain unbalanced cycles. This means that the edge is in a handcuff, and hence in $\mathrm{cl}_M(V)$. Now we conclude that each of the (at least) $a - \lambda$ edges is incident with at least one vertex counted by the variables $\gamma_1, \ldots, \gamma_y$. As the edges counted by a form a matching, we now see that $a - \lambda \leq \sum \gamma_i$. We conclude that

$$\begin{aligned} |N| = 2a + \sum \alpha_i + \sum \beta_i &\leq 2\sum \gamma_i + 2\lambda + \sum \alpha_i + \sum \beta_i \\ &\leq 2\sum \alpha_i + 6\sum \beta_i + 2\sum \gamma_i + 2\lambda. \end{aligned}$$

But (1) implies that $2\sum \alpha_i + 6\sum \beta_i + 2\sum \gamma_i \leq 12\lambda - 12$, so we are done. □

Remark 8.2. If we remove the constraint of 3-connectivity from Proposition 8.1, then no bound on the number of vertices in both $G[U]$ and $G[V]$ is possible. To see this, consider a cycle with an even number of edges, and alternately colour the edges blue and red. To each red edge attach a red clique, and to each blue edge attach a blue clique. The partition into blue and red edges is a 2-separation in the resulting graphic matroid, but the number of vertices incident with both blue and red edges is not bounded by any function of 2. It is possible to construct similar examples of frame matroids that are not graphic.

We will concentrate on two subclasses of frame matroids. *Bicircular* matroids are those that arise from empty linear classes. Thus every cycle is unbalanced. For any graph, G, we define $B(G)$ to be the bicircular matroid $M(G, \emptyset)$. Bicircular matroids can also be characterised as the transversal matroids represented by systems of the form $(A_1, \ldots, A_r)$, where each element of the ground set is in at most two of the sets $A_1, \ldots, A_r$.

Next we define gain-graphic matroids. Again, we let G be an undirected graph with edge set E and (possibly) loops and multiple edges. Define $A(G)$

to be

$$\{(e,u,v)\colon e \text{ is a non-loop edge joining vertices } u \text{ and } v\}$$
$$\cup \{(e,u,u)\colon e \text{ is a loop incident with the vertex } u\}.$$

A *gain function*, σ, takes $A(G)$ to a group H, and satisfies $\sigma(e,u,v) = \sigma(e,v,u)^{-1}$ for any non-loop edge e with end-vertices u and v. If $P = v_0e_0v_1e_1\cdots e_tv_{t+1}$ is a path of G, then the *gain value* of P is $\sigma(P) = \sigma(e_0,v_0,v_1)\cdots\sigma(e_t,v_t,v_{t+1})$. Now let $C = v_0e_0v_1e_1\cdots e_tv_{t+1}$ be a cycle of G, where $v_0 = v_{t+1}$, and the other vertices are pairwise distinct. Then $\sigma(C)$ is also defined to be $\sigma(e_0,v_0,v_1)\cdots\sigma(e_t,v_t,v_{t+1})$. Note that $\sigma(C)$ may depend on the choice of orientation of C, and if H is nonabelian, it may also depend on the choice of starting vertex. However, if $\sigma(C)$ is equal to the identity, then this equality will hold no matter which starting vertex and orientation we choose. We declare a cycle to be balanced exactly when $\sigma(C)$ is equal to the identity, and this gives rise to a linear class. If $\mathcal{B}$ is such a linear class, then $M(G,\mathcal{B})$ is an *H-gain-graphic* matroid. Gain-graphic matroids play an important role in the works by Kahn and Kung [11], and Geelen, Gerards, and Whittle [7]. We artificially close the classes of bicircular and H-gain-graphic matroids under the addition of matroid loops, in order to make the classes minor-closed.

Let u be a vertex of G, and let α be an element of H. The gain function $\sigma_{u,\alpha}$ is defined to be identical to σ on any edge not incident with u and on any loop. Furthermore $\sigma_{u,\alpha}(e,u,v) = \alpha\sigma(e,u,v)$ when e is a non-loop edge joining u to a vertex v, and in this case $\sigma_{u,\alpha}(e,v,u)$ is defined to be $\sigma(e,v,u)\alpha^{-1}$. The operation that produces $\sigma_{u,\alpha}$ from σ is called *switching*. Two gain functions that are related by switching have exactly the same balanced cycles [20, Lemma 5.2].

Proposition 8.3. *Let $\sigma\colon A(G) \to H$ be a gain function on the graph G, and let X be a subset of edges such that $G[X]$ is balanced. If u and v are distinct vertices of $G[X]$, and P and P' are paths in $G[X]$ from u to v, then $\sigma(P) = \sigma(P')$.*

Proof. Because $G[X]$ is balanced, we can repeatedly apply switching operations to produce a gain function that takes any edge in $G[X]$ to the identity of H [20, Lemma 5.3]. Therefore, after applying these switching operations, the gain value of any path from u to v is the identity of H. We can apply switching operations again to recover the original gain function, σ. We easily check that after applying these switching operations any two paths from u to v still have identical gain values. □

The next theorem treats bicircular matroids and gain-graphic matroids simultaneously, since the arguments are essentially identical.

Theorem 8.4. *The class of 3-connected bicircular matroids is computably pigeonhole. If H is a finite group, then the class of 3-connected H-gain-graphic matroids is computably pigeonhole.*

Proof. The succinct representation of a bicircular matroid, $M(G, \emptyset)$, is just a description of the graph, G. An H-gain-graphic matroid is described via a graph, and a labelling that assigns an element of H to each orientation of an edge. We assume that M is a 3-connected matroid with ground set E, and that M is either bicircular, or H-gain-graphic. Let G be the graph that represents M, so that G is unlabelled if M is bicircular, and labelled if M is H-gain-graphic. We can assume that G has no isolated vertices. Because M is 3-connected, this means that G is connected.

Let (U, V) be a partition of E such that $\lambda_M(U) \leq \lambda$ for some positive integer λ. Let N be the set of vertices that are in both $G[U]$ and $G[V]$, so that $|N| \leq 14\lambda - 12$ by Proposition 8.1.

Let X be an independent subset of U. We define the *signature* of X. This will contain a partition of the set of vertices in N that are also in $G[X]$. For each connected component, D, of $G[X]$, such that D contains vertices of N, we let the set of vertices in both N and D be a block of the partition. In addition, we record whether D is unbalanced or balanced. Furthermore, if M is H-gain-graphic and D is balanced, then for every pair of distinct vertices in both N and D, we record the gain value of a path in D that joins the pair. Proposition 8.3 tells us that this gain value is well-defined, and does not depend on the choice of path. To decide whether D is balanced, we let F be a spanning tree of D. We apply switching operations in such a way that each edge in F is labelled with the identity element of H. Now D is balanced if and only if every edge in D is labelled by the identity element in the new gain function [20, Lemma 5.3]. This procedure can clearly be accomplished in polynomial time, so it is clear that the signature of X can be computed in polynomial time.

Let X and X' be subsets of U. We declare X and X' to be equivalent under $\approx_U$ if both are dependent, or both are independent and they have identical signatures. Since we regard the number of elements in H as being fixed, Proposition 8.1 tells us that the number of $(\approx_U)$-classes is a function of λ, and does not depend on the choice of M, U, and V. Hence we can complete the proof of Theorem 8.4 by proving that $\approx_U$ refines $\sim_U$.

To this end, assume that X and X' are independent subsets of U, and that $X \approx_U X'$. Assume that $X' \cup Z$ is dependent for some $Z \subseteq V$. Let C' be a circuit of M contained in $X' \cup Z$, where we assume that C' contains edges from both X' and Z. We will replace edges in $C' \cap X'$ with edges from X and obtain another dependent set. This will establish that $X \cup Z$ is dependent, and we will be done.

For each connected component, D', of $G[X']$ such that D' contains edges of C', we do the following. Let $C'_1, \ldots, C'_t$ be subsets of X' such that $G[C'_1], \ldots, G[C'_t]$ are exactly the connected components of $G[C' \cap X']$ that are contained in D'. First assume that D' is unbalanced. Then there is a unbalanced connected component of $G[X]$, call it D, such that D and D' intersect N in exactly the same vertices. Therefore for each i, there is a set of edges, $C_i \subseteq X$, such that $G[C_i]$ is an unbalanced connected subgraph contained in D, and C_i contains all the vertices of N that are in $G[C_i']$.

Now assume that D' is balanced. Let D be the balanced connected component of $G[X]$ such that D and D' intersect N in exactly the same vertices, and furthermore if G is a gain-graph, then paths in D and D' between the same vertices of N have the same gain values. For each i, we let $C_i \subseteq X$ be chosen so that $G[C_i]$ is a balanced connected subgraph contained in D such that every vertex of C_i' in N is also in C_i.

Now we remove each C_i' from C', and replace it with the set C_i. We perform this operation for each connected component, D' of $G[X']$ that contains edges of C'. Let C be the set of edges that we obtain in this way, so that C is contained in $X \cup Z$. It is clear that $G[C]$ is connected. We will show that C is dependent in M, and this will complete the proof.

For any graph, Γ, let $\nu(\Gamma)$ be $|E(\Gamma)| - |V(\Gamma)|$. If Γ is connected, then $\nu(\Gamma) \geq -1$. If Γ is connected and contains exactly one cycle, then $\nu(\Gamma) = 0$. Let (L, R) be a partition of $E(\Gamma)$, and assume that γ vertices are incident with edges in both L and R. It is easy to confirm that

$$\nu(\Gamma) = \nu(\Gamma[L]) + \nu(\Gamma[R]) + \gamma. \tag{2}$$

If X is a subset of E, then X is dependent in M if and only if $G[X]$ contains a balanced cycle, or $\nu(G[X]) \geq 1$. Note that since each $G[C_i']$ and $G[C_i]$ is connected, $\nu(G[C_i']), \nu(G[C_i]) \geq -1$. If $G[C_i']$ contains no cycle, then $\nu(G[C_i]) \geq \nu(G[C_i']) = -1$. If $G[C_i']$ contains a cycle, then that cycle must be unbalanced, for we have assumed that the circuit C' is not contained in X'. As C_i' is independent, it follows that C_i' contains exactly one cycle, so $\nu(G[C_i']) = 0$. In this case, our choice of a substitute component $G[C_i]$ also contains an unbalanced cycle, so $\nu(G[C_i]) \geq \nu(G[C_i']) = 0$. In any case, $G[C_i]$ shares at least as many vertices with $G[Z]$ as $G[C_i']$ does. Now it follows from (2) that removing C_i' from $G[C']$ and replacing it with C_i produces a subgraph with at least the same value of ν. In other words, $\nu(G[C]) \geq \nu(G[C'])$. If $\nu(G[C]) \geq 1$, then $G[C]$ is a connected subgraph containing at least two cycles, so C is dependent, and we are done.

Therefore we assume that $\nu(G[C]) < 1$. As $G[C']$ contains at least one cycle, we see that $\nu(G[C]) = \nu(G[C']) = 0$. Therefore $G[C']$ is a cycle, and it must be a balanced cycle. Now each component $G[C_i']$ is a path. If any component $G[C_i]$ contains an unbalanced cycle, then $\nu(G[C])$ will be greater than $\nu(G[C'])$, a contradiction. Therefore each component D must be a balanced component. This means that each $G[C_i]$ contains a path joining the end-vertices of $G[C_i']$, and these paths have the same gain value. Now we can easily see that $G[C]$ also contains a balanced cycle, and again we conclude that C is dependent. □

Corollary 8.5. *Let $\mathcal{M}$ be the class of bicircular or H-gain-graphic matroids (with H a finite group). Let ψ be any sentence in MS_0. There is a fixed-parameter tractable algorithm for testing whether matroids in $\mathcal{M}$ satisfy ψ, where the parameter is branch-width.*

Proof. This will follow immediately from [6, Theorem 6.7] and Theorem 8.4 if we show that the succinct representations of bicircular and H-gain-graphic matroids are minor-compatible. We rely on [20, Corollary 5.5] and [21, Theorem 2.5]. Let M be a bicircular or H-gain-graphic matroid corresponding to the graph G, and let e be an edge of G. Then $M\backslash e$ is bicircular or H-gain-graphic, and corresponds to $G\backslash e$. (In the case that M is H-gain-graphic, the edge-labels in $G\backslash e$ are inherited from G.)

Contraction is somewhat more difficult. If e is a non-loop, then we first perform a switching (in the H-gain-graphic case) so that the gain-value on e is the identity. We then simply contract e from G. The resulting graph represents M/e. Now assume e is a loop of G incident with the vertex u. If e is a balanced loop, we simply delete e, so now assume that e is an unbalanced loop. In the H-gain-graphic case, this implies that H is non-trivial. We obtain the graph G' by deleting u and replacing each non-loop edge, e', incident with u with a loop incident with the other end-vertex of e'. In the H-gain-graphic case, the loop e' is labelled with any non-identity element. Any other loops of G that are incident with u are added as matroid loops after contracting e (remembering that we closed the classes of bicircular and H-gain-graphic matroids under the addition of such elements).

It is clear that the operations of deletion and contraction can be performed in polynomial time, so the classes of bicircular and H-gain-graphic matroids have minor-compatible succinct representations as desired. □

Now the proof of Theorem 1.2 is complete, by Corollary 8.5.

Remark 8.6. Hliněný has shown [8, p. 348] that his work provides an alternative proof of Courcelle's Theorem. We can provide a simple new proof by relying on Corollary 8.5, as we now briefly explain.

Let ψ be a sentence in the monadic second-order logic, MS_2 of graphs. This means that we can quantify over variables representing vertices, edges, sets of vertices and set of edges. We have binary predicates for set membership, and also an incidence predicate, which allows us to express that an edge is incident with a vertex. We need to show that there is a fixed-parameter tractable algorithm for testing ψ in graphs, with respect to the parameter of tree-width.

Let G be a graph, and let G° be the graph obtained from G by adding two loops at every vertex. We need to interpret ψ as a sentence about bicircular matroids of the form G°. We let $\mathrm{Vert}(X_i)$ be the MS_0 formula stating that X_i is a 2-element circuit. Similarly, we let $\mathrm{Edge}(X_i)$ be a formula expressing that X_i is a singleton set not contained in a 2-element circuit. Now we make the following interpretations in ψ: if Q is a quantifier and v is a vertex variable, we replace Qv with $QX_v\,\mathrm{Vert}(X_v)$. If e is an edge variable, we replace Qe with $QX_e\,\mathrm{Edge}(X_e)$. If V is a variable representing a set of vertices, we replace QV with

$$QX\forall X_1(\mathrm{Sing}(X_1)\wedge X_1\subseteq X)\rightarrow \exists X_2(X_1\subseteq X_2\wedge X_2\subseteq X\wedge \mathrm{Vert}(X_2)).$$

Similarly, if E is a variable representing a set of edges, then we replace QE with $QX\forall X_1(\mathrm{Sing}(X_1) \wedge X_1 \subseteq X \rightarrow \mathrm{Edge}(X_1))$. Finally, we replace any occurrence of the predicate stating that e is incident with v with an MS_0 formula saying that there is a 3-element circuit that contains X_e and one of the elements in X_v. We let ψ' be the sentence we obtain by making these substitutions. It is clear that a graph, G, satisfies ψ if and only if $B(G^\circ)$ satisfies ψ'. Therefore Corollary 8.5 implies that there is a fixed-parameter tractable algorithm for testing whether ψ' holds in matroids of the form $B(G^\circ)$, with respect to the parameter of branch-width.

To find the branch-width of a graph with edge set E, we consider a ternary tree, T, and a bijection from E to the leaves of T. If (U, V) is a partition of E displayed by an edge, e, of T, then we count the vertices incident with edges in both U and V. This gives us the *width* of e, and the maximum width of an edge of T is the width of the decomposition. The lowest width across all such decompositions is the branch-width of the graph. It is not difficult to see that the branch-width of the matroid $B(G^\circ)$ is bounded by a function of the branch-width of the graph G, and similarly the branch-width of G is bounded by a function of the branch-width of $B(G^\circ)$. But exactly the same relation holds between the branch-width and the tree-width of G [17, (5.1)]. Now it follows that there is a fixed-parameter tractable algorithm for testing whether ψ holds in graphs, where the parameter is tree-width. This proves Courcelle's Theorem [4].

When H is not finite, the class of H-gain-graphic matroids is not even pigeonhole, as we now show. First we require the following proposition.

Proposition 8.7. *Let H be an infinite group, and let m and n be positive integers. There are disjoint subsets $A, B \subseteq H$ such that $|A| = m$, $|B| = n$, and $\{ab\colon (a, b) \in A \times B\}$ is disjoint from $A \cup B$ and has cardinality mn.*

Proof. Assume that $m = 1$. Choose B, an arbitrary subset of n elements that does not include the identity. The cancellation rule implies the result if we let A be a singleton set containing an element not in $B \cup \{b_1 b_2^{-1}\colon b_1, b_2 \in B\}$. The result similarly holds if $n = 1$. Now we let m and n be chosen so that $m + n$ is as small as possible with respect to the proposition failing. Let A' and B be disjoint subsets such that $|A'| = m - 1$, $|B| = n$, and $\{ab\colon (a, b) \in A' \times B\}$ has cardinality $(m-1)n$ and is disjoint from A' and B. We choose an element x not in $A' \cup B$ that does not belong to $\{ab^{-1}\colon a \in A, b \in B\}$, nor to $\{b_1 b_2^{-1}\colon b_1, b_2 \in B\}$, nor to $\{ab_1 b_2^{-1}\colon a \in A, b_1, b_2 \in B\}$. Now we simply let A be $A \cup \{x\}$. $\square$

Proposition 8.8. *Let H be an infinite group. There are rank-3 H-gain-graphic matroids with arbitrarily high decomposition-width. Hence the class of H-gain-graphic matroids is not pigeonhole.*

Proof. Assume otherwise, and let K be an integer such that $\mathrm{dw}(M) \leq K$ whenever M is a rank-3 H-gain-graphic matroid.

Znám proved that if a bipartite graph with n vertices in each side of its bipartition has more than $(d-1)^{1/d}n^{2-1/d}+n(d-1)/2$ edges, then it has a subgraph isomorphic to $K_{d,d}$ [22]. Choose an integer d such that $d^2>K$. Choose the integer p so that

$$\frac{1}{2}p^2>(d-1)^{1/d}p^{2-1/d}+\frac{1}{2}p(d-1).$$

Finally, choose the integer q such that $q-p\geq q/2\geq p$ and

$$\frac{1}{3}(q^2+2q)-p(2q-p+2)>(d-1)^{1/d}(q-p)^{2-1/d}+\frac{1}{2}(q-p)(d-1).$$

Using Proposition 8.7, we choose disjoint subsets $A=\{a_1,\ldots,a_q\}$ and $B=\{b_1,\ldots,b_q\}$ of H such that $a_ib_j\neq a_pb_q$ whenever $(i,j)\neq(p,q)$. Let AB be $\{a_ib_j\colon 1\leq i,j\leq q\}$. We also assume that AB is disjoint from $A\cup B$. Let G be a graph on vertex set $\{v_1,v_2,v_3\}$, where there are q parallel edges between v_1 and v_2 and between v_2 and v_3, and q^2 parallel edges between v_1 and v_3. We let σ be the gain function applying the elements in A to the q arcs from v_1 to v_2, the elements in B to the arcs from v_2 to v_3, and the elements in AB to those arcs from v_1 to v_3. We identify these group elements with the ground set of the H-gain-graphic matroid $M=M(G,\sigma)$. Therefore M is a rank-3 matroid with ground set $A\cup B\cup AB$. Its non-spanning circuits are the 3-element subsets of A, B, or AB, along with any set of the form $\{a_i,b_j,a_ib_j\}$.

Let (T,φ) be a decomposition of M with the property that if U is any displayed set, then $\sim_U$ has at most K equivalence classes. As in the proof of Lemma 4.1, we let e be an edge of T such that each of the displayed sets, U_e and V_e, contains at least $|E(M)|/3=(q^2+2q)/3$ elements. We construct a complete bipartite graph with vertex set $A\cup B$ and edge set AB, where a_ib_j joins a_i to b_j. We colour a vertex or edge red if it belongs to U_e, and blue otherwise. Without loss of generality, we will assume that at least $q/2\geq p$ vertices in A are red.

Assume that B contains at least p blue vertices. We choose p such vertices, and p red vertices from A, and let G' be the graph induced by these $2p$ vertices. There are p^2 edges in G'. Assume that at least $p^2/2$ of them are red (the case that at least $p^2/2$ of them are blue is almost identical). Our choice of p means that G' contains a subgraph isomorphic to $K_{d,d}$ consisting of red edges. Thus there are elements $a_{i_1},\ldots,a_{i_d}\in A\cap U_e$ and $b_{j_1},\ldots,b_{j_d}\in B\cap V_e$ such that every element $a_{i_p}b_{j_q}$ is in U_e. For $(l,k)\neq(p,q)$, we see that $\{a_{i_l},a_{i_l}b_{j_k}\}$ is not equivalent to $\{a_{i_p},a_{i_p}b_{j_q}\}$, since $\{a_{i_l},a_{i_l}b_{j_k},b_{j_k}\}$ is a circuit of M, and $\{a_{i_p},a_{i_p}b_{j_q},b_{j_k}\}$ is a basis. Therefore $\sim_{U_e}$ has at least $d^2>K$ equivalence classes, and we have a contradiction. We must now assume that B contains fewer than p blue vertices, and hence at least $q-p\geq q/2$ red vertices. Thus a symmetrical argument shows that A contains fewer than p blue vertices.

We choose $q-p$ red vertices from each of A and B, and let G'' be the subgraph induced by these vertices. Let g stand for the number of blue edges in G''. The number of edges not in G'' is equal to $q^2 - (q-p)^2 = 2pq - p^2$. As there are g blue edges in G'', at most $2pq - p^2$ blue edges not in G'', and fewer than $2p$ blue vertices, it follows that $|V_e| < g + 2pq - p^2 + 2p$. Since $(q^2 + 2q)/3 \leq |V_e|$, we deduce that

$$\frac{1}{3}(q^2 + 2q) - p(2q - p + 2) < g.$$

Our choice of q now means that G'' has a subgraph isomorphic to $K_{d,d}$ consisting of blue edges. Thus we have elements $a_{i_1}, \ldots, a_{i_d} \in A \cap U_e$ and $b_{j_1}, \ldots, b_{j_d} \in B \cap U_e$ such that $a_{i_p} b_{j_q}$ is in V_e for each p and q. For $(l,k) \neq (p,q)$, we see that $\{a_{i_l}, b_{j_k}, a_{i_l} b_{j_k}\}$ is a circuit of M, while $\{a_{i_p}, b_{j_q}, a_{i_l} b_{j_k}\}$ is a basis. This implies there are at least d^2 equivalence classes under $\sim_{U_e}$, so we again have a contradiction. $\square$

9. Open problems

We have proved that the class of lattice path matroids is pigeonhole, but we have been unable to prove that it is strongly pigeonhole. Nevertheless, we believe this to be the case.

Conjecture 9.1. *The class of lattice path matroids is computably pigeonhole.*

The classes of fundamental transversal matroids and lattice path matroids are both closed under duality ([16, Proposition 11.2.28] and [2, Theorem 3.5]). Thus they belong to the intersection of transversal and cotransversal matroids. We suspect that Theorem 6.3 (and Conjecture 9.1) exemplify a more general result.

Conjecture 9.2. *The class of matroids that are both transversal and cotransversal is strongly pigeonhole.*

Despite the existence of examples as in Remark 8.2, we firmly believe the next conjecture.

Conjecture 9.3. *The class of bicircular matroids is computably pigeonhole. Let H be a finite group. The class of H-gain-graphic matroids is computably pigeonhole.*

10. Acknowledgements

We thank Geoff Whittle for several important conversations, and Yves de Cornulier for providing a proof of Proposition 8.7. Funk and Mayhew were supported by a Rutherford Discovery Fellowship, managed by Royal Society Te Apārangi.

References

[1] H. L. Abbott, D. Hanson, and N. Sauer. Intersection theorems for systems of sets. *J. Combinatorial Theory Ser. A* **12** (1972), 381–389.

[2] J. Bonin, A. de Mier, and M. Noy. Lattice path matroids: enumerative aspects and Tutte polynomials. *J. Combin. Theory Ser. A* **104** (2003), no. 1, 63–94.

[3] J. E. Bonin, J. P. S. Kung, and A. de Mier. Characterizations of transversal and fundamental transversal matroids. *Electron. J. Combin.* **18** (2011), no. 1, Paper 106, 16.

[4] B. Courcelle. The monadic second-order logic of graphs. I. Recognizable sets of finite graphs. *Inform. and Comput.* **85** (1990), no. 1, 12–75.

[5] P. Erdős and R. Rado. Intersection theorems for systems of sets. *J. London Math. Soc.* **35** (1960), 85–90.

[6] D. Funk, D. Mayhew, and M. Newman. Tree automata and pigeonhole classes of matroids – I. *arXiv e-prints* (2019), arXiv:1910.04360.

[7] J. Geelen, B. Gerards, and G. Whittle. Structure in minor-closed classes of matroids. In *Surveys in combinatorics 2013*, volume 409 of *London Math. Soc. Lecture Note Ser.*, pp. 327–362. Cambridge Univ. Press, Cambridge (2013).

[8] P. Hliněný. Branch-width, parse trees, and monadic second-order logic for matroids. *J. Combin. Theory Ser. B* **96** (2006), no. 3, 325–351.

[9] A. W. Ingleton and M. J. Piff. Gammoids and transversal matroids. *J. Combin. Theory Ser. B* **15** (1973), no. 1, 51–68.

[10] M. M. Jose and D. Mayhew. Well-quasi-ordering in lattice path matroids. *arXiv e-prints* (2018), arXiv:1806.10260.

[11] J. Kahn and J. P. S. Kung. Varieties of combinatorial geometries. *Trans. Amer. Math. Soc.* **271** (1982), no. 2, 485–499.

[12] D. Kőnig. Graphs and matrices. *Mat. és Fiz. Lapok* **38** (1931), 116–119.

[13] D. Král. Decomposition width of matroids. *Discrete Appl. Math.* **160** (2012), no. 6, 913–923.

[14] D. Mayhew, M. Newman, and G. Whittle. On excluded minors for real-representability. *J. Combin. Theory Ser. B* **99** (2009), no. 4, 685–689.

[15] D. Mayhew, M. Newman, and G. Whittle. Yes, the "missing axiom" of matroid theory is lost forever. *Trans. Amer. Math. Soc.* **370** (2018), no. 8, 5907–5929.

[16] J. Oxley. *Matroid theory.* Oxford University Press, New York, second edition (2011).

[17] N. Robertson and P. D. Seymour. Graph minors. X. Obstructions to tree-decomposition. *J. Combin. Theory Ser. B* **52** (1991), no. 2, 153–190.

[18] Y. Strozecki. *Enumeration complexity and matroid decomposition.* Ph.D. thesis, Université Paris Diderot (2010).

[19] Y. Strozecki. Monadic second-order model-checking on decomposable matroids. *Discrete Appl. Math.* **159** (2011), no. 10, 1022–1039.

[20] T. Zaslavsky. Biased graphs. I. Bias, balance, and gains. *J. Combin. Theory Ser. B* **47** (1989), no. 1, 32–52.

[21] T. Zaslavsky. Biased graphs. II. The three matroids. *J. Combin. Theory Ser. B* **51** (1991), no. 1, 46–72.

[22] Š. Znám. On a combinatorical problem of K. Zarankiewicz. *Colloq. Math.* **11** (1963), 81–84.

Faculty of Science and Technology, Douglas College, Vancouver, Canada
Email address: funkd@douglascollege.ca

School of Mathematics and Statistics, Victoria University of Wellington, New Zealand
Email address: dillon.mayhew@vuw.ac.nz

Department of Mathematics and Statistics, University of Ottawa, Ottawa, Canada
Email address: mnewman@uottawa.ca