Tree automata and pigeonhole classes of matroids: II

Daryl Funk[a] Dillon Mayhew[b] Mike Newman[c]

Submitted: Nov 6, 2022; Accepted: May 24, 2023; Published: Jul 14, 2023
© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

Let ψ be a sentence in the counting monadic second-order logic of matroids and let $\mathbb{F}$ be a finite field. Hliněný's Theorem says that we can test whether $\mathbb{F}$-representable matroids satisfy ψ using an algorithm that is fixed-parameter tractable with respect to branch-width. In a previous paper we proved there is a similar fixed-parameter tractable algorithm that can test the members of any efficiently pigeonhole class. In this sequel we apply results from the first paper and thereby extend Hliněný's Theorem to the classes of fundamental transversal matroids, lattice path matroids, bicircular matroids, and H-gain-graphic matroids, when H is a finite group. As a consequence, we can obtain a new proof of Courcelle's Theorem.

Mathematics Subject Classifications: 05B35, 68Q27, 68Q45, 68R99

1 Introduction

In the first paper of the series [8], we proved an extension of Hliněný's Theorem [10]. This theorem concerns the *counting monadic second-order logic* for matroids, CMS_0. In this language we have variables $X_1, X_2, X_3, \ldots$ representing subsets of the ground set of a matroid. We have a binary predicate $X_i \subseteq X_j$, which allows us to say when one subset is contained in another. We also have a unary independence predicate $\mathsf{Ind}(X_i)$, which returns the value true when the input is an independent subset. We further have predicates that allow us to say that a set has cardinality p modulo q, where p and q are positive integers. Let ψ be a sentence in CMS_0. 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 $\mathbb{F}$. 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

[a] Department of Mathematics, Douglas College, British Columbia, Canada. (funkd@douglascollege.ca).
[b] School of Mathematics and Statistics, Te Herenga Waka, Victoria University of Wellington, Wellington, New Zealand. (dillon.mayhew@vuw.ac.nz).
[c] Department of Mathematics and Statistics, University of Ottawa, Ontario, Canada. (mnewman@uottawa.ca).

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 [8] extends Hliněný's Theorem.

Theorem 1. *Let $\mathcal{M}$ be an efficiently pigeonhole class of matroids. Let ψ be a sentence in CMS_0. We can test whether matroids in $\mathcal{M}$ satisfy ψ using an algorithm that is fixed-parameter tractable with respect to branch-width.*

This sequel paper exploits Theorem 1 and related ideas to show that there are fixed-parameter tractable algorithms for testing CMS_0 sentences in other natural classes of matroids, beyond finite-field representable matroids. In particular, our main theorem shows that Hliněný's Theorem can be extended as follows.

Theorem 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 CMS_0. We can test whether matroids in $\mathcal{M}$ satisfy ψ using an algorithm that is fixed-parameter tractable with respect to branch-width.

We now explain efficiently 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 $E(M)$. We define an equivalence relation on the subsets of U. 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 an appropriately chosen subcubic 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 [14] and by Strozecki [18, 19], in the sense that the decomposition-width of a class is bounded under one definition if and only if it is bounded under the other.

A class of matroids with bounded decomposition-width must have bounded branch-width [8, Corollary 2.8]. The converse does not hold (Lemma 16). 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 ([8, Corollary 5.3]). The class of $\mathbb{F}$-representable matroids forms a pigeonhole class if and only if $\mathbb{F}$ is finite (Theorem 18 and Proposition 19). Fundamental transversal matroids (Theorem 22) and lattice path matroids also form pigeonhole classes (Corollary 31).

A stronger property arises quite naturally. Imagine that $\mathcal{M}$ is a class of matroids and that M is an arbitrary matroid in $\mathcal{M}$. Assume that whenever $U \subseteq E(M)$ and $\lambda_M(U) \leqslant k$, then there are at most $\pi(k)$ equivalence classes under $\sim_U$, where $\pi(k)$ is a value depending only on k. (Recall that $\lambda_M(U)$ is the connectivity value of U and is defined to be $r_M(U) + r_M(V) - r(M)$.) In this case we say that $\mathcal{M}$ is *strongly pigeonhole* (Definition 9). Every strongly pigeonhole class is pigeonhole [8, Proposition 2.11], but the converse does not hold [8, Remark 2.12]. 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 18). We do not know if any of the other classes in Theorem 2 are strongly pigeonhole, but we certainly believe this to be the case (Conjectures 39 and 41). In fact, we make the broad conjecture that the class of matroids that are transversal and cotransversal is a strongly pigeonhole class (Conjecture 40).

Theorem 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 *efficiently pigeonhole* (see Definition 14 for the formal statement). Any efficiently pigeonhole class is also strongly pigeonhole. The class of matroids representable over a finite field (Theorem 18) is efficiently pigeonhole, and this leads to a proof of Hliněný's Theorem. The class of fundamental transversal matroids is also efficiently pigeonhole (Theorem 22).

In [8, Theorem 6.11] we also proved that there is a fixed-parameter-tractable algorithm for testing ψ in members of $\mathcal{M}$ as long as the 3-connected members of $\mathcal{M}$ form an efficiently pigeonhole class and we can efficiently construct descriptions of minors. This result was motivated by the fact that we do not know if bicircular matroids or H-gain-graphic matroids (H finite) form efficiently pigeonhole classes. (We conjecture this is the case in Conjecture 41.) We have been able to show that the 3-connected bicircular matroids and the 3-connected H-gain-graphic matroids form efficiently pigeonhole classes (Theorem 34). This is then enough to establish cases (iii) and (iv) in Theorem 2.

Knowing that we have efficient model-checking for bicircular matroids gives us a new, and quite simple, proof of Courcelle's Theorem (Remark 36), 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 17). Thus matroids representable over a given infinite field form a non-pigeonhole class (Proposition 19). The class of transversal matroids is not pigeonhole, (Proposition 20) and nor is the class of gammoids (Remark 21). A different argument shows that the class of H-gain-graphic matroids is not pigeonhole when H is infinite (Proposition 38).

For an introduction to monadic second-order logic and its application to finite struc-

tures, see [6] or [7, Chapter 13]. Oxley provides our reference for the basic concepts of matroid theory [16]. We have noted that if M is a matroid and (U, V) is a partition of $E(M)$, then the *connectivity value* of U is $\lambda_M(U) = r_M(U) + r_M(V) - r(M)$. A *$k$-separation* is a partition, (U, V), of the ground set such that $|U|, |V| \geqslant k$, and $\lambda_M(U) < k$. A matroid is *n-connected* if it has no k-separations with $k < n$. A *cyclic flat* of the matroid M is a flat F such that the restriction $M|F$ has no coloops. When E is a finite set and $\mathcal{I}$ is a collection of subsets of E, we refer to the pair $(E, \mathcal{I})$ as a *set-system*. We say that the members of $\mathcal{I}$ are *independent sets*.

2 Tree automata

Definition 3. 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 (so that T cannot have exactly two vertices). 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 4. 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 *transition rules*, δ_0 and δ_2, are partial functions from Σ and $\Sigma \times Q \times Q$ respectively, into 2^Q, the power set of Q.

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 $r(v)$ is $\delta_0(\sigma(v))$ if this is defined, and is otherwise the empty set.

(ii) if v has left child v_L and right child v_R, then

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

as long as the images in this union are all defined: if they are not then we set $r(v)$ to be empty.

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

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 be a subset of E. We construct a $\Sigma \cup (\Sigma \times \{0, 1\})$-tree

which we denote $\mathrm{enc}(T, \sigma, \varphi, Y)$. The characters applied to the leaves of this tree will encode the subset Y. If v is a non-leaf vertex of T, then it receives the label $\sigma(v)$ in $\mathrm{enc}(T, \sigma, \varphi, Y)$. However, if v is a leaf, then it receives a label $(\sigma(v), s)$, where $s = 1$ if $\varphi^{-1}(v)$ is in Y and otherwise $s = 0$.

Definition 5. Let Σ be a finite set and let A be a tree automaton with $\Sigma \cup (\Sigma \times \{0, 1\})$ 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:

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

So $M(A, T, \sigma, \varphi)$ is the set-system consisting of all subsets of the leaves that are accepted by the automaton.

Now let Σ be a finite set, and let A be a tree automaton with alphabet $\Sigma \cup (\Sigma \times \{0, 1\})$. 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 $(T_M, \sigma_M, \varphi_M)$ is a *parse tree* for M (relative to the automaton A).

Note that if $(T_M, \sigma_M, \varphi_M)$ 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 $(\sigma_M(v), 1)$ if $\varphi_M^{-1}(v)$ is in Y and with $(\sigma_M(v), 0)$ otherwise. By then running A on the resulting tree, and testing to see if it accepts, we can check whether or not Y is in $\mathcal{I}$. This idea is central to the proofs of Hliněný's Theorem and of Theorem 1.

3 Pigeonhole classes

This section states our central definitions: decomposition-width, pigeonhole classes, strongly pigeonhole classes, and efficiently pigeonhole classes.

Definition 6. 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 *subcubic 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 subcubic 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 [14] 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. A class of matroids with bounded rank also has bounded branch-width. In [8, 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 16). This motivates the following definition.

Definition 7. 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 $\mathrm{bw}(M) \leqslant \lambda$ implies $\mathrm{dw}(M) \leqslant \rho(\lambda)$, 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 [8, Corollary 5.3].

Proposition 8. *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 9. 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 $\lambda_M(U) \leqslant \lambda$, there are at most $\pi(\lambda)$ equivalence classes under $\sim_U$. In this case we say that $\mathcal{M}$ is *strongly pigeonhole*.

In [8, Proposition 2.11], we give the easy proof that any class with the strong pigeonhole property also has the pigeonhole property. On the other hand, the class of rank-2 matroids is pigeonhole without being strongly pigeonhole (see [8, Remark 2.12].)

Proposition 10. *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 $\lambda_M(U) \leqslant \lambda$, for some positive integer λ. Declare subsets $X, X' \subseteq U$ to be equivalent if:

(i) $|X|, |X'| > r_M(U)$,

(ii) $r_M(U) - \lambda < |X| = |X'| \leqslant r_M(U)$, or

(iii) $|X|, |X'| \leqslant r_M(U) - \lambda$.

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 now need only consider the case that $|X|, |X'| \leqslant r_M(U) - \lambda$. 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 $|Z| \leqslant r_M(E(M) - U)$. Therefore

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

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

Theorem 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 11. 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, along with a polynomial-time Turing Machine. We write $\Delta(M)$ to indicate any string in the image of $M \in \mathcal{M}$. The Turing machine must, when given any input $(\Delta(M), X)$ where $M \in \mathcal{M}$ and X is a subset of $E(M)$, correctly answer the question "Is X independent in M?".

If the Turing Machine operates in time bounded by at most $p(n)$ on any input of length n, where p is some polynomial, then it follows 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.

Proposition 12. *Let Δ be a succinct representation of $\mathcal{M}$, a class of matroids. For each positive integer λ let $\mathcal{M}_\lambda$ be $\{M \in \mathcal{M} \colon \mathrm{bw}(M) \leqslant \lambda\}$. Let f be a function from positive integers to positive integers. Assume there is an algorithm which when given a positive integer λ will produce a tree automaton A_λ in time bounded by $f(\lambda)$. Assume also there is an algorithm which when given $(\Delta(M), \lambda)$ for $M \in \mathcal{M}_\lambda$ will produce a parse tree for M relative to A_λ in time bounded by $f(\lambda)|\Delta(M)|^c$, where c is a constant. Let ψ be a sentence in CMS_0. We can test whether matroids in $\mathcal{M}$ satisfy ψ using an algorithm that is fixed-parameter tractable with respect to branch-width. Furthermore, $\mathcal{M}$ is pigeonhole.*

Proof. Assume we are given the input $(\Delta(M), \lambda)$, where M is in $\mathcal{M}_\lambda$. We first construct A_λ in time bounded by $f(\lambda)$. Assume that $(T_M, \sigma_M, \varphi_M)$ is the parse tree that we construct for $M \in \mathcal{M}_\lambda$ relative to A_λ. This takes time bounded by $f(\lambda)|\Delta(M)|^c$. Now [8, Lemma 4.7] tells us that there is a tree automaton $A_{\lambda,\psi}$ such that $A_{\lambda,\psi}$ accepts $(T_M, \sigma_M, \varphi_M)$ if and only M satisfies ψ. Moreover, the proof of this lemma is constructive, and it tells us how to build $A_{\lambda,\psi}$, given A_λ. The time taken to build $A_{\lambda,\psi}$ depends only λ and ψ (which we regard as fixed). Let this time be equal to $g_\psi(\lambda)$.

Applying $A_{\lambda,\psi}$ to $(T_M, \sigma_M, \varphi_M)$ takes time bounded by $h_\psi(\lambda)|\Delta(M)|$, for some value $h_\psi(\lambda)$ that depends only on ψ and λ. Now the total process of testing M to see if it satisfies ψ takes time bounded by

$$f(\lambda)(|\Delta(M)|^c + 1) + g_\psi(\lambda) + h_\psi(\lambda)|\Delta(M)|.$$

This establishes the existence of the fixed-parameter tractable algorithm.

The existence of A_λ and $(T_M, \sigma_M, \varphi_M)$ for any $M \in \mathcal{M}_\lambda$ means that $\mathcal{M}_\lambda$ is *automatic* (using the language of [8]). Because it is automatic, it has bounded decomposition-width ([8, Lemma 5.1]). So for any positive integer λ, the class $\mathcal{M}_\lambda$ has bounded decomposition-width. This is exactly what it means for $\mathcal{M}$ to be pigeonhole. □

Definition 13. 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 $\Delta(M/X\backslash Y)$.

In order to construct parse trees for automata 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 14. Let $\mathcal{M}$ be a class of matroids with a succinct representation Δ. Assume there is a Turing Machine, a constant c, and a function π from positive integers to positive integers such that the machine takes as input any tuple $(\Delta(M), U, X, X', \lambda)$, where M is in $\mathcal{M}$, $U \subseteq E(M)$ satisfies $\lambda_M(U) \leqslant \lambda$, and X and X' are subsets of U. Assume also that in time bounded by $O(\pi(\lambda)|E(M)|^c)$, the machine computes an equivalence relation, $\approx_U$, on the subsets of U, so that it accepts $(\Delta(M), U, X, X', \lambda)$ if and only if $X \approx_U X'$. Furthermore we assume that

(i) $X \approx_U X'$ implies $X \sim_U X'$, and

(ii) the number of equivalence classes under $\approx_U$ is at most $\pi(\lambda)$.

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

Remark 15. The proof of Theorem 1 essentially follows from Proposition 12 and the observation that given $\Delta(M)$ we can construct a branch-decomposition for M in polynomial time ([8, Proposition 6.3]). We can then convert this decomposition tree into a parse tree. The automaton which processes this tree can also be constructed in time bounded by $O(\lambda^c)$. Theorem 1 follows[1].

Clearly an efficiently pigeonhole class of matroids is also strongly pigeonhole. In Theorem 22, we will prove that the class of fundamental transversal matroids is efficiently pigeonhole. Statement (i) of Theorem 2 will then immediately follow, by an application of Theorem 1.

4 Non-pigeonhole classes

Next we develop some tools for proving negative results. We want to certify that certain classes are not pigeonhole. Recall that a matroid with rank r is *sparse paving* if every

[1]Note that the proof of Theorem 1 requires a definition of efficiently pigeonhole classes that is very slightly stronger than the one presented in [8].

circuit has cardinality either r or $r+1$ and when C and C' are distinct circuits of size r then $|C \cap C'| < r-1$. Let G be a simple graph with edge set $\{e_1, \ldots, e_m\}$ and vertex set $\{v_1, \ldots, v_n\}$, where $n \geqslant 3$. 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 16. *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 k be an arbitrary integer greater than K, and let N be an integer such that

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

and $\mathcal{M}$ contains a matroid, M, isomorphic to $m(K_N)$. By relabelling, we assume that the ground set of M is $\{v_1, \ldots, v_N\} \cup \{e_{ij} \colon 1 \leqslant i < j \leqslant N\}$ 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

$$\frac{1}{3}|E(M)| = \frac{1}{3}\left(N + \binom{N}{2}\right) \geqslant 7k^2 + 2k$$

elements. Let G be a complete graph with vertex set $\{v_1, \ldots, v_N\}$ and edge set $\{e_{ij} \colon 1 \leqslant i < j \leqslant N\}$, 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. Then we can assume without loss of generality 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

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

A result of Abbott, Hanson, and Sauer [1] implies that the subgraph induced by red edges that join two blue vertices 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 (relative to 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 17. *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 16. □

5 Representable matroids

The next result is not surprising, and is implicitly utilised in the work of Hliněný [10] and Král [14].

Theorem 18. *Let $\mathbb{F}$ be a finite field. The class of $\mathbb{F}$-representable matroids is efficiently 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 $\lambda_M(U) \leqslant \lambda$. Grassman's identity tells us that the rank of $\langle U\rangle \cap \langle V\rangle$ is equal to $r(U) + r(V) - r(M) \leqslant \lambda$. 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 $X \approx_U X'$ is true 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 [10] follows immediately from Theorems 1 and 18. We note that proofs of Hliněný's Theorem can also be derived from the works by Král [14] and Strozecki [19].

Proposition 19. *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 17 and [15, Lemma 2.1]. □

6 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]). In this case B is a basis consisting of points located at the vertices of the simplex. From this characterisation it is easy to see that the dual of a fundamental transversal matroid is also fundamental.

In this, and subsequent sections, we will show that three subclasses of transversal matroids are pigeonhole: fundamental transversal matroids (Theorem 22), lattice path matroids (Theorem 30), and bicircular matroids (Theorem 34). To start with, we prove that we cannot extend these results to the entire class of transversal matroids.

Proposition 20. *The class of transversal matroids is not pigeonhole.*

Proof. By Proposition 8, 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* [11]. 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 17. □

Remark 21. From Proposition 20 we see that any class of matroids containing transversal matroids is not pigeonhole. In particular, the class of gammoids is not pigeonhole.

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, W, of G such that $|W| = |X \cap A|$ and each edge in W joins a vertex in $X \cap A$ to a vertex in $B - X$. In this case we say that W *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. 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'. The transversal matroid on the ground set $A \cup B$ represented by this bipartite graph is equal to $M[G]$.

Theorem 22. *The class of fundamental transversal matroids is efficiently pigeonhole.*

Proof. We consider the succinct representation of fundamental transversal matroids that involves representing such a matroid with a bipartite graph. Let $M = M[G]$ be a fundamental transversal matroid, where $A \cup B$ is a bipartition of the bipartite graph G, and B is a basis of M. Let (U, V) be a partition of $A \cup B$, and assume that $\lambda_M(U) \leqslant \lambda$. Our goal will be to construct an equivalence relation $\approx$ on the subsets of U such that $\approx$ satisfies the conditions in Definition 14.

Let H be the subgraph of G induced by edges that join vertices in $A \cap U$ to vertices in $B \cap V$, and vertices in $A \cap V$ to vertices in $B \cap U$.

Claim 23. *Any matching of H contains at most λ edges.*

Proof. Let W 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 W. Therefore $|A_U| + |A_V| = |W|$. If we restrict W 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

$$\lambda \geqslant r(U) + r(V) - r(M) \geqslant |B \cap U| + |A_U| + |B \cap V| + |A_V| - |B| = |W|. \qquad \square$$

We can find a maximum matching of H, using one of a number of polynomial-time algorithms. It follows from Kőnig's Theorem [13] that H contains a vertex cover, S, such that $|S| \leqslant \lambda$. 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 W be a matching that certifies its independence. We will construct a *signature*, $\mathcal{C}(X, W)$. Signatures of subsets of V will be defined symmetrically at the same time, so in fact we let $\{P, Q\}$ be $\{U, V\}$, and we let X be an independent subset of P, with W a matching certifying the independence of X. Recall that this means that $|W| = |X \cap A|$ and each edge of W joins a vertex in $X \cap A$ to a vertex in $B - X$.

The signature $\mathcal{C}(X, W)$ 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 $\mathcal{C}(X, W)$ 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 W.

(ii) A subset $Z \subseteq A \cap Q \cap S$ is in $\mathcal{S}_2$ if and only if there is a matching W' satisfying $W \subseteq W'$ and $|W' - W| = |Z|$, where each edge in $W' - W$ 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 W to a vertex in $(B \cap Q) - S$. Note that any vertex in S_3 is in X, since it is in A and incident with an edge of W.

(iv) S_4 is the set of vertices in $B \cap Q \cap S$ that are incident with an edge in W.

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 W, which is drawn with heavy lines. Vertices in the sets S_1, S_3, and S_4 are marked. The sets in the family $\mathcal{S}_2$ are the empty set and the singleton set containing the vertex marked S_2.

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

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

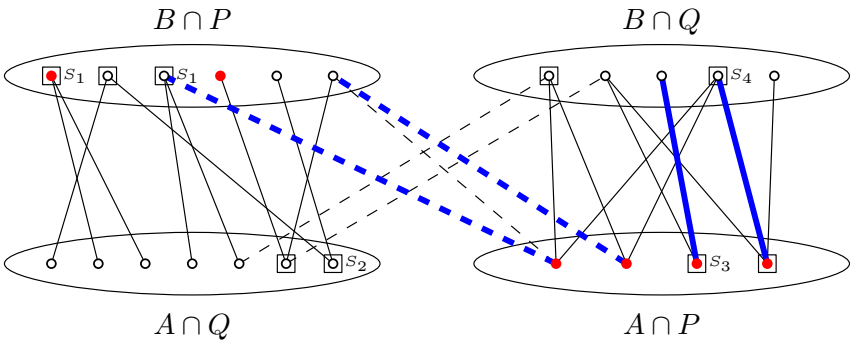

Figure 1: Defining a signature.

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 $\mathcal{C}(X) = \mathcal{C}(X')$. Thus independent sets X and X' are equivalent under $\approx_U$ if they have exactly the same signatures.

We can choose a signature by choosing three subsets of S, and a family of subsets of S. Recalling that $|S| \leqslant \lambda$, we see that the number of signatures is at most $(2^\lambda)^3 \cdot 2^{2^\lambda} = 2^{3\lambda+2^\lambda}$. Therefore the number of $(\approx_U)$-classes is no more than

$$\pi(\lambda) = 2^{2^{3\lambda+2^\lambda}} + 1. \tag{1}$$

Now we have established that condition (ii) in Definition 14 holds, where the function $\pi(\lambda)$ is defined by (1). To complete the proof that the class of fundamental transversal matroids is efficiently pigeonhole, we must establish that we can compute the equivalence relation in time bounded by $O(\pi(\lambda)|E(M)|^c)$ and that $\approx_U$ refines $\sim_U$. Let us first consider the task of computing $\approx_U$. Recall that $\{P, Q\} = \{U, V\}$.

Claim 24. *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 W, certifying the independence of X, such that $\mathcal{C}(X, W) = (S_1, \mathcal{S}_2, S_3, S_4)$ 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 $S - S_1$. 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 $S - S_4$.

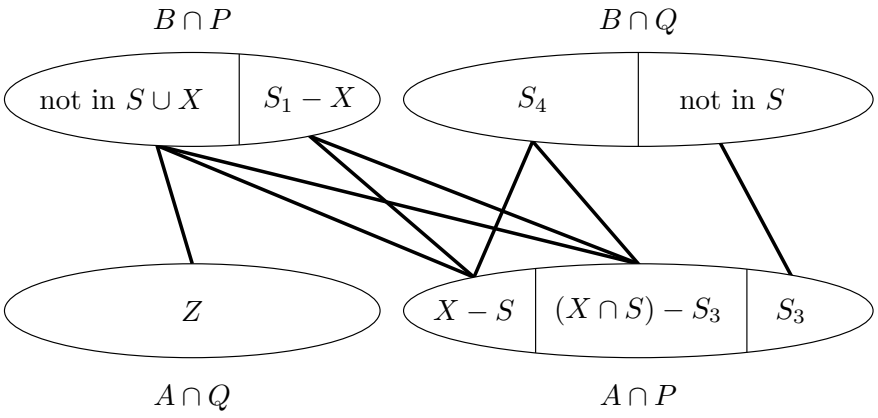

Figure 2: The construction of G'.

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 $(B \cap Q) - S$. Obviously G' can be constructed in time bounded by $O(|E(M)|^c)$, for some constant c.

24.1 The following statements are equivalent:

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

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

Proof. Assume (i) holds. Then $|W| = |X \cap A|$, and every vertex in $X \cap A$ is incident with an edge of W. Furthermore, every edge of W joins a vertex in $X \cap A$ to a vertex in $B - X$. Since Z is in $\mathcal{S}_2$, we can let W' be a matching such that $W \subseteq W'$, $|W' - W| = |Z|$, and each edge of $W' - W$ joins a vertex in Z to a vertex in $(B \cap P) - (S \cup X)$. We will show that every edge of W' is in G', and hence W' is a matching of G'.

Let $ab \in W'$ be an edge joining $a \in A$ to $b \in B$. Certainly every edge of $W' - W$ is an edge of G'. Therefore we will assume that ab is in W. Then a is in X, and hence in P. Note that b is not in X, for no edge of W joins two vertices in X. Assume that a is in $X - S$. If b is in P, then it is either in $(B \cap P) - (S \cup X)$, or it is in S. In the latter case, b is in S_1 by the definition of S_1. In either case, ab is an edge of G'. Now assume b is in Q. Then b must be in S, or else ab is an edge of H not incident with a vertex in S. In this case b is in S_4 by definition, so ab is an edge of G'.

Next assume a is in $(X \cap S) - S_3$. Assume b is in P. Then either b is in $(B \cap P) - (S \cup X)$, or it is in $S_1 - X$. In either case ab is an edge of G'. Now assume b is in Q. If b is not in S, then the definition of S_3 implies $a \in S_3$, since a is in S. But this contradicts the assumption $a \in (X \cap S) - S_3$. Therefore b is in S. This means b is in S_4, so ab is an edge of G'.

Finally, we 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'. Now we have shown that W' is a matching of

G'. Every vertex in $(X \cup Z) \cap A$ is incident with an edge in W', and the same statement is true for vertices in $(S_1 - X) \cup S_4$, by the definitions of S_1 and S_4. Therefore (ii) holds.

Now assume (ii) holds. Let W' 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 W'. Let W be the set of edges in W' 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, the matching W certifies the independence of X. Every vertex in $S_1 - X$ is incident with an edge of W, and no vertex of $(B \cap P \cap S) - S_1$ is (since these vertices are not in G). Therefore in $\mathcal{C}(X, W)$, the first entry is S_1, as desired. Every edge of $W' - W$ joins a vertex in Z to a vertex in $(B \cap P) - (S \cup X)$, so W' certifies that Z belongs to the second entry of $\mathcal{C}(X, W)$. Because S_3 is a subset of X we see that any vertex in S_3 is matched by W to a vertex in $(B \cap Q) - S$. Furthermore no vertex in $(X \cap S \cap A) - S_3$ is, by the construction of G'. Therefore the third entry of $\mathcal{C}(X, W)$ is equal to S_3. Finally, every vertex in S_4 is incident with an edge in W, and no vertex of $(B \cap Q \cap S) - S_4$ is (since these vertices are not in G'). Therefore $\mathcal{C}(X, W) = (S_1, \mathcal{S}_2, S_3, S_4)$, where Z is in $\mathcal{S}_2$, so (i) holds. □

Now we complete the proof of Claim 24. To test whether W 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. It is easy to see that we can use alternating-path methods to test whether there is a complete 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 24.1 justifies the correctness of this algorithm. □

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 $\mathcal{C}(X')$. By using Claim 24, we can accomplish this in time bounded by $O(\pi(\lambda)|E(M)|^c)$, for some constant c, where $\pi(\lambda)$ is the function in (1).

Now our final task in the proof of Theorem 22 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. 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 $\mathcal{C}(Y)$. Note that $S_1, T_4 \subseteq B \cap U$ and $S_4, T_1 \subseteq B \cap V$, while $\mathcal{S}_2$ is a family of subsets of $A \cap V$ and $\mathcal{T}_2$ is a family of subsets of $A \cap U$. We also have $S_3 \subseteq A \cap U$ and $T_3 \subseteq A \cap V$. 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) $S_4 \cap T_1 = \emptyset$.

We will prove that $X \cup Y$ is independent in M if and only if we can find signatures in $\mathcal{C}(X)$ and $\mathcal{C}(Y)$ that are compatible. This task is completed in Claim 25, and its converse (Claim 26). From this it will easily follow that $\approx_U$ refines $\sim_U$, and that therefore the class of fundamental transversal matroids is efficiently pigeonhole.

Claim 25. *Let $X \subseteq U$ and $Y \subseteq V$ be independent subsets of M. If $X \cup Y$ is independent in M 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 W be a matching certifying that $X \cup Y$ is independent. Thus $|W| = |X \cup Y|$, and every edge of W joins a vertex in $(X \cup Y) \cap A$ to a vertex in $B - (X \cup Y)$. Let W_X and W_Y be the subsets of W consisting of edges incident with vertices in X (respectively Y). Then W_X certifies the independence of X, and W_Y certifies the independence of Y. We assert that the signatures $\mathcal{C}(X, W_X)$ and $\mathcal{C}(Y, W_Y)$ are compatible.

Let $\mathcal{C}(X, W_X)$ be $(S_1, \mathcal{S}_2, S_3, S_4)$ and let $\mathcal{C}(Y, W_Y)$ 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 W_X. On other hand, T_4 is the set of vertices in $B \cap U \cap S$ incident with an edge of W_Y. No edge in W_Y is incident with an edge in W_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 incident with an edge of W_X 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 W_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 W_Y to a vertex in $(B \cap U) - S$. Let W' be the union of W_X along with the set of edges in W_Y that are incident with a vertex in T_3. Clearly W' is a matching as it is a subset of W. Also, $W_X \subseteq W'$ and $|W' - W_X| = |T_3|$. Each edge in $W' - W_X$ 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 W join vertices in $(X \cup Y) \cap A$ to vertices in $B - (X \cup Y)$. Therefore each edge in $W' - W_X$ joins a vertex of T_3 to one in $(B \cap U) - (S \cup X)$. We have established that T_3 is contained in $\mathcal{S}_2$. The symmetrical argument shows that S_3 is in $\mathcal{T}_2$. Therefore $\mathcal{C}(X, W_X)$ and $\mathcal{C}(Y, W_Y)$ are compatible, as we claimed. □

Claim 26. *Let $X \subseteq U$ and $Y \subseteq V$ be independent subsets of M. 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.*

Proof. We assume that $\mathcal{C}(X, W_X) = (S_1, \mathcal{S}_2, S_3, S_4)$ and $\mathcal{C}(Y, W_Y) = (T_1, \mathcal{T}_2, T_3, T_4)$ are compatible signatures. Recall that S_1 and T_4 are subsets of $B \cap U$ and that S_4 and T_1 are subsets of $B \cap V$. Furthermore T_3 is a subset of $A \cap V$ and $\mathcal{S}_2$ is a family of subsets of $A \cap V$. Finally, S_3 is a subset of $A \cap U$ and $\mathcal{T}_2$ is a family of subsets of $A \cap U$. 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 W_X to vertices in $(B \cap V) - S$. Let W''_X be the subset of W_X containing edges that are incident with vertices in S_3. The compatibility of the signatures means that S_3 is in $\mathcal{T}_2$. Hence there is a matching, W'_Y, such that $W_Y \subseteq W'_Y$, $|W'_Y - W_Y| = |S_3|$, and each edge of $W'_Y - W_Y$ joins a vertex in S_3 to one in $(B \cap V) - (S \cup Y)$.

Similarly, we let W''_Y be the subset of W_Y containing edges that are incident with vertices in T_3. Thus each edge in W''_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 W'_X be a matching such that $W_X \subseteq W'_X$, $|W'_X - W_X| = |T_3|$, and each edge of $W'_X - W_X$ joins a vertex in T_3 to a vertex in $(B \cap U) - (S \cup X)$. We now make the definition

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

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

26.1 *W is a matching.*

Proof. Note that $W'_X - W''_X$ and $W'_Y - W''_Y$ are certainly matchings. So if W is not matching then there is a vertex w, and distinct edges $wx \in W'_X - W''_X$ and $wy \in W'_Y - W''_Y$.

We first assume that w is in A. Assume also that w is in U. Because Y is a subset of V, no edge of W_Y is incident with a vertex in $A \cap U$. Therefore wy is in $W'_Y - W_Y$. 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 $W'_X - W_X$ is incident with a vertex in $A \cap U$, so wx is not in $W'_X - W_X$. Therefore it is in W_X, so wx is an edge of W_X that is incident with a vertex in S_3. But this means that wx is in W''_X, contradicting $wx \in W'_X - W''_X$. If w is in V, then we reach the similar contradiction that wy is in W''_Y. Therefore we must now assume that w is in B.

We assume that w is in $B \cap U$. No edge of $W'_Y - W_Y$ is incident with a vertex in $B \cap U$. Therefore wy is in W_Y, and it follows that 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, wy joins an edge of $A \cap V \cap S$ to a vertex in $(B \cap U) - S$. This implies that y is in T_3. Now wy belongs to W''_Y, and we have a contradiction to $wy \in W'_Y - W''_Y$. Thus w is in S.

If wx is in $W'_X - W_X$, 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 W_X. Now w is in S and is incident with an edge of W_X, meaning that it is in S_1. Furthermore, the edge wy is in W_Y, and this means that w is in T_4. Thus $S_1 \cap T_4 \neq \emptyset$, and we have a contradiction to the fact that $\mathcal{C}(X, W_X)$ and $\mathcal{C}(Y, W_Y)$ are compatible.

If w is in V, then we reach the symmetric contradiction that either wx is in W''_X, or w is in $S_4 \cap T_1$. This completes the proof that W is a matching. □

Now we must demonstrate that the matching W certifies the independence of $X \cup Y$. First we show that every edge of W joins a vertex in $(X \cup Y) \cap A$ to a vertex in $B - (X \cup Y)$. To this end, we let ab be an edge of W, where a is in A and b is in B. Without loss of generality we can assume that ab is in $W'_X - W''_X$. Then a is either in X or in T_3, which is a subset of Y. Therefore a is in $(X \cup Y) \cap A$.

We demonstrate that b is not in $X \cup Y$. If ab is in $W'_X - W_X$, 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 $W'_X - W_X$, so it is in W_X. Each edge of W_X joins a vertex of $A \cap X$ to a vertex of $B - X$, so b is not in X. Assume that b is in Y, so that it belongs to $B \cap V$.

Since $X \subseteq U$, it follows that a is in $A \cap U$. 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 W_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 W''_X, a contradiction. Therefore b is in S. As b is in $B \cap Y \cap S$, it follows that it is in T_1. But the edge ab also certifies that b is in S_4. Therefore $S_4 \cap T_1 \neq \emptyset$, and we have a contradiction to the fact that $\mathcal{C}(X, W_X)$ and $\mathcal{C}(Y, W_Y)$ are compatible. We have shown that b is not in $X \cup Y$. Hence every edge of W joins a vertex of $A \cap (X \cup Y)$ to a vertex in $B - (X \cup Y)$.

In the final step we must show that every vertex of $(X \cup Y) \cap A$ is incident with an edge in W. Let a be a vertex in $(X \cup Y) \cap A$. Without loss of generality, we will assume that a is in $A \cap X$. Then a is certainly incident with an edge of W_X. If it is not incident with an edge of W, then it is not incident with an edge in $W'_X - W''_X$. It follows that in this case, a is incident with an edge of W''_X. This implies a is in S_3. But in this case w is incident with an edge of $W'_Y - W_Y$, and hence with an edge of W. Thus any vertex of $A \cap (X \cup Y)$ is incident with an edge in W, so W certifies the independence of $X \cup Y$, exactly as we desired. □

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 25 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 26 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 22 is complete. □

Case (i) in Theorem 2 follows immediately from Theorem 22 and Theorem 1.

Remark 27. Theorem 22 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 21 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]. Every lattice path matroid is transversal. We describe an algorithm that constructs a parse tree for a given lattice path matroid (Theorem 30). When combined with Proposition 12 this shows that there is a fixed-parameter tractable algorithm for testing CMS_0 sentences in lattice path matroids. It also shows that the class is pigeonhole.

A lattice path matroid is represented by a pair of strings made from the alphabet $\{E, N\}$. Let P be such a string, so that $P = p_1 p_2 \cdots p_s$, where each p_i is equal to either E or N. When i is in $\{1, \ldots, s\}$, we let $n_i(P)$ stand for the number of N-characters in $\{p_1, \ldots, p_i\}$. We also let $N(P)$ stand for $\{i \in \{1, \ldots, s\} \colon p_i = N\}$. Thus $n_s(P) = |N(P)|$. If P and Q are strings of s characters from the alphabet $\{E, N\}$ then we write $P \preccurlyeq Q$ to mean that $n_i(P) \leqslant n_i(Q)$ for each $i \in \{1, \ldots, s\}$.

Now assume that P and Q each contain r copies of N and m copies of E. Any such string can be identified with a path in the integer lattice from $(0,0)$ to (m,r) that uses only North and East steps. If $P \preccurlyeq Q$ then the path P never goes above Q. In this case, an *intermediate string*, L, is composed of r copies of N and m copies of E and satisfies $P \preccurlyeq L \preccurlyeq Q$. Note that P and Q are both intermediate strings.

Let P and Q be strings composed of r copies of N and m copies of E such that $P \preccurlyeq Q$. The *lattice path matroid* $M[P,Q]$ has $\{1,\ldots,m+r\}$ as its ground set. The family of bases of $M[P,Q]$ is $\{N(L)\colon L \text{ is an intermediate string}\}$.

As we have mentioned, every minor of a lattice path matroid is a lattice path matroid. We now give an explicit description of such minors, following [4, p. 707]. Imagine that $M = M[P,Q]$ is a lattice path matroid, where P and Q contain r copies of N and m copies of E. Let P be $p_1p_2\cdots p_{m+r}$ and let Q be $q_1q_2\cdots q_{m+r}$. Let i be in $\{1,\ldots,m+r\}$. Assume that $n_i(P) = n_i(Q)$. Then i is a coloop of M if $p_i = q_i = N$ and i is a loop if $p_i = q_i = E$. In either of these cases, both $M\backslash i$ and M/i are equal to

$$M[p_1p_2\cdots p_{i-1}p_{i+1}\cdots p_{m+r}, q_1q_2\cdots q_{i-1}q_{i+1}\cdots q_{m+r}]$$

after we relabel the elements $i+1, i+2, \ldots, m+r$ in $M\backslash i$ or M/i as $i, i+1, \ldots, m+r-1$.

Now we assume that neither of these scenarios applies. Let j be the largest integer in $\{1,\ldots,i\}$ such that $p_j = E$, and let k be the least integer in $\{i,\ldots,m+r\}$ such that $q_k = E$. Then

$$M\backslash i = M[p_1p_2\cdots p_{j-1}p_{j+1}\cdots p_{m+r}, q_1q_2\cdots q_{k-1}q_{k+1}\cdots q_{m+r}],$$

where we apply exactly the same relabelling as before. Next let j be the least integer in $\{i,\ldots,m+r\}$ such that $p_j = N$, and let k be the largest integer in $\{1,\ldots,i\}$ such that $q_k = N$. In this case

$$M/i = M[p_1p_2\cdots p_{j-1}p_{j+1}\cdots p_{m+r}, q_1q_2\cdots q_{k-1}q_{k+1}\cdots q_{m+r}].$$

Proposition 28. *Let $M = M[P,Q]$ be a lattice path matroid, where P and Q contain r copies of N and m copies of E. If $n_i(Q) - n_i(P) \geqslant t$ for some $i \in \{1,\ldots,m+r\}$, then M contains a minor isomorphic to $U_{t,2t}$.*

Proof. Let P be $p_1\cdots p_{m+r}$ and let Q be $q_1\cdots q_{m+r}$. Assume that the result fails, so that M has no $U_{t,2t}$-minor. Furthermore assume that M has been chosen subject to this failure so that $m+r$ is as small as possible.

If $p_1 = N$, then $q_1 = N$ because otherwise $P \preccurlyeq Q$ fails. In this case $M/1$ is isomorphic to $M[p_2\cdots p_{m+r}, q_2\cdots q_{m+r}]$. But now $q_2,\ldots,q_i$ contains at least t more copies of N than $p_2,\ldots,p_i$. Since $M/1$ has no $U_{t,2t}$-minor, the minimality of M has been contradicted. Therefore $p_1 = E$. Symmetrically, if $q_{m+r} = N$, then $p_{m+r} = N$ or else $n_{m+r-1}(P) > n_{m+r-1}(Q)$. In this case $M/(m+r)$ is isomorphic to $M[p_1\cdots p_{m+r-1}, q_1\cdots q_{m+r-1}]$. But if $i < m+r$ then the first i characters of $p_1\cdots p_{m+r-1}$ are the same as the first i characters of $p_1\cdots p_{m+r}$, and the same applies to $q_1\cdots q_{m+r-1}$. If $i = m+r$, then $q_1\cdots q_{m+r-1}$ contains at least t more N-characters than $p_1\cdots p_{m+r-1}$. In either case we see that $M/(m+r)$ is a smaller counterexample, so $q_{m+r} = E$.

Let $k \in \{1, \ldots, m+r\}$ be the least integer such that $q_k = E$, and let $Q' = q_1 \cdots q_{k-1} q_{k+1} \cdots q_{m+r}$. Let P' be $p_2 \cdots p_{m+r}$. Thus $M\backslash 1$ is isomorphic to $M[P', Q']$. Now $n_{i-1}(P') = n_i(P)$. If $k \leqslant i$, then $n_{i-1}(Q') = n_i(Q)$, and in this case $n_{i-1}(Q') - n_{i-1}(P') \geqslant t$, so $M\backslash 1$ provides a smaller counterexample. Therefore $k > i$, which means that $q_1, \ldots, q_i$ are all equal to N.

Assume there is at least one N-character in $p_1, \ldots, p_i$. Let $j \in \{1, \ldots, i\}$ be the least integer such that $p_j = N$. Let $Q'' = q_2 \cdots q_{m+r}$ and let P'' be $p_1 \cdots p_{j-1} p_{j+1} \cdots p_{m+r}$. Thus $M/1$ is isomorphic to $M[P'', Q'']$. In this case

$$n_{i-1}(Q'') - n_{i-1}(P'') = (i-1) - (n_i(P) - 1) = i - n_i(P) = n_i(Q) - n_i(P) \geqslant t$$

so $M/1$ is a smaller counterexample. Hence $p_1, \ldots, p_i$ are all equal to E.

Note that $i \geqslant t$. Assume that $i > t$. Then

$$n_{i-1}(Q') - n_{i-1}(P') = (i-1) - 0 \geqslant t,$$

which implies that $M\backslash 1$ is a smaller counterexample. We conclude that $i = t$.

Assume that there is an integer $j \in \{i+1, \ldots, m+r\}$ such that $p_j = E$, and let j be the largest such integer. We have noted that $q_{m+r} = E$. Now we define P' to be $p_1 \cdots p_{j-1} p_{j+1} \cdots p_{m+r}$ and Q' to be $q_1 \cdots q_{m+r-1}$. Thus $M\backslash(m+r)$ is isomorphic to $M[P', Q']$. But $n_i(Q') = n_i(Q) = t$ and $n_i(P') = n_i(P) = 0$, so we have a smaller counterexample. Therefore $p_{i+1}, \ldots, p_{m+r}$ are all equal to N.

Note that $m + r \geqslant 2t$. Assume that $m + r > 2t$. Since $p_1, \ldots, p_i$ are E-characters and $p_{i+1}, \ldots, p_{m+r}$ are N-characters, and $i = t$, this means that P and Q contain more than t copies of N. Let k be the largest integer such that $q_k = N$. Since Q contains more than $i = t$ copies of N, we see that $k > i$. Define Q'' to be $q_1 \cdots q_{k-1} q_{k+1} \cdots q_{m+r}$ and define P'' to be $p_1 \cdots p_{m+r-1}$. Then $M/(m+r)$ is isomorphic to $M[P'', Q'']$. Again we see that $n_i(Q'') = n_i(Q) = t$ and $n_i(P'') = n_i(P) = 0$. Now we have a smaller counterexample, so $m + r = 2t$.

We have shown that P consists of t copies of E followed by t copies of N. The first t characters in Q are N, so the following t characters must be E. It is now very easy to see that M is isomorphic to $U_{t,2t}$, so we have a final contradiction. □

Corollary 29. *Let $M = M[P, Q]$ be a lattice path matroid, where P and Q contain r copies of N and m copies of E. Assume that $\mathrm{bw}(M) \leqslant \lambda$. Then $n_i(Q) - n_i(P) \leqslant \frac{1}{2}(3\lambda - 3)$ for each $i \in \{1, \ldots, m+r\}$.*

Proof. Let t be equal to $n_i(Q) - n_i(P)$ for an arbitrary choice of i in $\{1, \ldots, m+r\}$. Then Proposition 28 implies that M has an $U_{t,2t}$-minor. The branch-width of this minor is $\lceil 2t/3 \rceil + 1$ by [16, Exercise 14.2.5]. Now [16, Proposition 14.2.3] implies that $\lceil 2t/3 \rceil + 1 \leqslant \mathrm{bw}(M) \leqslant \lambda$. The result follows. □

In the next result, we aim to apply Proposition 12 to the class of lattice path matroids. Lattice path matroids are represented succinctly by pairs of strings.

Theorem 30. *For each positive integer λ there is a tree automaton A_λ that can be constructed in time bounded by $f(\lambda)$, for some function f on the positive integers, so that when we are given (P, Q, λ), where $M = M[P, Q]$ is a lattice path matroid with branch-width at most λ, we can construct a parse tree for M relative to A_λ in time bounded by $f(\lambda)|E(M)|^c$, where c is a constant.*

Proof. Let p be $\lfloor \frac{1}{2}(3\lambda - 3) \rfloor$. Let $\mathcal{P}$ be the power set of $\{1, 2, \ldots, p\}$. Let Σ be the set that contains the empty function ϵ, as well as all functions from $\mathcal{P} \times \{0, 1\}$ or from $\{0, 1\} \times \{0, 1\}$ into $\mathcal{P}$. Then A_λ has $\Sigma \cup \{(\epsilon, 0), (\epsilon, 1)\}$ as its alphabet. The state space of A_λ is $\mathcal{P}$, and all states are accepting except for the empty set.

Assume that P and Q contain r copies of N and m copies of E. Thus the ground set of M is $\{1, \ldots, m + r\}$. Let $P = p_1 p_2 \cdots p_{m+r}$ and let $Q = q_1 q_2 \cdots q_{m+r}$. Assume that $\mathrm{bw}(M) \leqslant \lambda$. We construct the tree T_M as shown in Figure 3. The bijection φ_M takes the element $i \in \{1, \ldots, m + r\}$ to the leaf v_i.

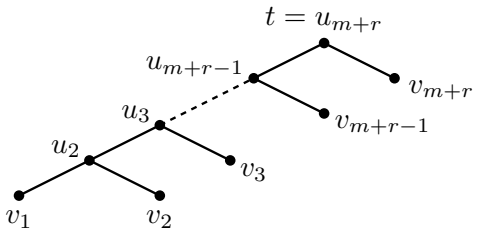

Figure 3: The parse tree for a lattice path matroid.

We imagine that Y is a subset of $\{1, 2, \ldots, m+r\}$. The idea behind the operation of A_λ is that it will apply a member of $\mathcal{P}$ to each of the internal nodes of T_M. The set applied to u_i will contain all possible values $|N(L_i)| - |N(p_1 p_2 \cdots p_i)|$, where L_i ranges over all strings of i characters satisfying $p_1 p_2 \cdots p_i \preccurlyeq L_i \preccurlyeq q_1 q_2 \cdots q_i$ and $Y \cap \{1, 2, \ldots, i\} \subseteq N(L_i)$. Note that Corollary 29 implies that such a value can be at most p and that Y is independent in M if and only if the set applied to the root t is non-empty.

To implement this idea, we describe σ_M, which applies a function to each node of T_M. We say that σ_M applies the empty function ϵ to any leaf. Next we let $g \in \Sigma$ be the function applied by σ_M to u_2. The domain of g is $\{0, 1\} \times \{0, 1\}$. The output of g depends on the characters in $p_1 p_2$ and $q_1 q_2$. In the following table, the columns are labelled by the array

$$\begin{matrix} q_1 q_2 \\ p_1 p_2 \end{matrix}$$

and each row shows the output of g on the members of $\{0, 1\} \times \{0, 1\}$.

	$(0,0)$	$(0,1)$	$(1,0)$	$(1,1)$
EE EE	$\{0\}$	$\emptyset$	$\emptyset$	$\emptyset$
EN EE	$\{0,1\}$	$\{1\}$	$\emptyset$	$\emptyset$
NE EE	$\{0,1\}$	$\{1\}$	$\{1\}$	$\emptyset$
NN EE	$\{0,1,2\}$	$\{1,2\}$	$\{1,2\}$	$\{2\}$
EN EN	$\{0\}$	$\{0\}$	$\emptyset$	$\emptyset$
NE EN	$\{0\}$	$\{0\}$	$\{0\}$	$\emptyset$
NN EN	$\{0,1\}$	$\{0,1\}$	$\{0,1\}$	$\{1\}$
NE NE	$\{0\}$	$\emptyset$	$\{0\}$	$\emptyset$
NN NE	$\{0,1\}$	$\{1\}$	$\{0,1\}$	$\{1\}$
NN NN	$\{0\}$	$\{0\}$	$\{0\}$	$\{0\}$

Next we let g be the function applied to u_i by σ_M, where i is in

$$\{3,\dots,m+r\}.$$

The domain of g will be $\mathcal{P}\times\{0,1\}$ and the output of g will be in $\mathcal{P}$. Assume that the input to g is (A,s). First consider the case that $p_i = N$. Then

$$g(A,0) = (A\cup\{a-1\colon a\in A\})\cap\{0,1,\dots,n_i(Q)-n_i(P)\}$$

and

$$g(A,1) = A\cap\{0,1,\dots,n_i(Q)-n_i(P)\}.$$

On the other hand, if $p_i = E$ then

$$g(A,0) = (A\cup\{a+1\colon a\in A\})\cap\{0,1,\dots,n_i(Q)-n_i(P)\}$$

and

$$g(A,1) = \{a+1\colon a\in A\}\cap\{0,1,\dots,n_i(Q)-n_i(P)\}.$$

Next we must describe the transition functions δ_0 and δ_2. The former takes each pair (ϵ,s) to $\{s\}$, where s is 0 or 1. If g is a function from $\{0,1\}\times\{0,1\}$ or $\mathcal{P}\times\{0,1\}$ to $\mathcal{P}$ then $\delta_2(g,j,k)$ is $g(j,k)$ whenever (j,k) is in the domain of g. This completes the description of the automaton

$$A_\lambda = (\Sigma\cup\{(\epsilon,0),(\epsilon,1)\},\mathcal{P},\mathcal{P}-\{\emptyset\},\delta_0,\delta_2).$$

Clearly A_λ can be constructed in time bounded by $f(\lambda)$, where f is some function from the positive integers to the positive integers.

To complete the proof of Theorem 30, we must show that $(T_M, \sigma_M, \varphi_M)$ is a parse tree relative to the automaton A_λ. We can accomplish this by proving our claim that the set applied to u_i by A_λ during its run on $\mathrm{enc}(T_M, \sigma_M, \varphi_M, Y)$ is the set $\{|N(L_i)| - |N(p_1p_2\cdots p_i)|\}$, where L_i ranges over $\{E, N\}^i$ subject to $p_1p_2\cdots p_i \preccurlyeq L_i \preccurlyeq q_1q_2\cdots q_i$ and $Y \cap \{1, 2, \ldots, i\} \subseteq N(L_i)$. When $i = 2$, this is simply a matter of checking the table above, so we inductively assume the claim is true for $i-1$. Let A be the set applied to u_{i-1}. If $p_i = N$ and i is in Y, then the final character in L_i must be N, so the set applied to u_i should be exactly the numbers in A, as long as those numbers are not greater than $n_i(Q) - n_i(P)$. But this set is exactly $g(A, 1)$, where g is the function applied to u_i by the labelling σ_M. So in this case, the set applied to u_i is exactly as we claimed.

If i is not in Y, then the final character in L_i could be either E or N, as long as the constraint $p_1p_2\cdots p_i \preccurlyeq L_i \preccurlyeq q_1q_2\cdots q_i$ is satisfied. Thus the set applied to u_i should contain all numbers of the form $a-1$ or a, where a ranges over the members of A, as long as these numbers are in $\{0, 1, \ldots, n_i(Q) - n_i(P)\}$. But this is exactly the output $g(A, 0)$. The case when $p_i = E$ yields to exactly the same sort of analysis, so $(T_M, \sigma_M, \varphi_M)$ is a parse tree relative to A_λ. It is clear that $(T_M, \sigma_M, \varphi_M)$ can be constructed in time bounded by $f(\lambda)(m + r)^c$ for some constant c, so the proof is complete. □

By applying Proposition 12 to the previous result, we see that case (ii) in Theorem 2 is proved. Furthermore, we can deduce the following result.

Corollary 31. *The class of lattice path matroids is pigeonhole.*

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 minimal path that joins 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 32. *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 $\lambda_M(U) \leqslant \lambda$. 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. We can assume that G has no isolated vertices. It then follows from the 3-connectivity of M that G is connected. 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 $G[V]$ contains 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 $\alpha_i \geqslant 3$ 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 $|\gamma_i| \geqslant 3$, $|N| = 2x + \sum \gamma_i + \sum \delta_i$, and $r_M(V) = n_V - (x + y)$. Because G is connected, $r(M) \geqslant n - 1$, and $r(M) = n - 1$ if and only if G is balanced. Now we observe that

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

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

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

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

$$6(\lambda - 1) \geqslant \sum \alpha_i + 3 \sum \beta_i + \sum \gamma_i + 3 \sum \delta_i. \tag{2}$$

The edges counted by a form a matching. Therefore they are an independent set in M. As $r_M(U)+r_M(V)-r(M) \leqslant \lambda$, 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 $a-\lambda \leqslant \sum \gamma_i$. We conclude that

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

But (2) implies that $2\sum \alpha_i + 6\sum \beta_i + 2\sum \gamma_i \leqslant 12\lambda - 12$, and the result follows. □

Remark 33. If we remove the constraint of 3-connectivity from Proposition 32, then no bound on the number of vertices in both $G[U]$ and $G[V]$ is possible. To see this, let $c_0, \ldots, c_{2n-1}$ be vertices in a cycle of the 2-connected graph G. Assume that $G - \{c_i, c_j\}$ is disconnected for any $i \neq j$. Define $\mathcal{B}$ to be the family of cycles that contain all of the vertices $c_0, \ldots, c_{2n-1}$. It is easy to verify that $\mathcal{B}$ is a linear class. For any i (modulo $2n$) let C_i be the set of edges contained in a path from c_i to c_{i+1} containing no other vertex in $c_0, \ldots, c_{2n-1}$. Then

$$(C_0 \cup C_2 \cup \cdots \cup C_{2n-2}, C_1 \cup C_3 \cup \cdots \cup C_{2n-1})$$

is a 1-separation of $M(G, \mathcal{B})$, but obviously there is no bound on the number of vertices incident with edges in both sides of this separation.

We will concentrate on two subclasses of frame matroids. *Bicircular* matroids are those frame matroids arising from linear classes that contain only loops. Thus every cycle with more than one edge is unbalanced. For any graph, G, we define $B(G)$ to be the bicircular matroid $M(G, \emptyset)$. Thus every bicircular matroid is equal to $B(G) \oplus U_{0,t}$ for some graph G and some integer t. 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 $W = v_0e_0v_1e_1\cdots e_tv_{t+1}$ is a walk of G, then the *gain-value* of W is $\sigma(W) = \sigma(e_0,v_0,v_1)\cdots\sigma(e_t,v_t,v_{t+1})$. 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. 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 [12], and Geelen, Gerards, and Whittle [9].

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 loop and on any edge not incident with u. 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].

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

Theorem 34. *The class of* 3*-connected bicircular matroids is efficiently pigeonhole. If H is a finite group, then the class of* 3*-connected H-gain-graphic matroids is efficiently pigeonhole.*

Proof. Let M be a 3-connected matroid that is either bicircular or H-gain-graphic. If M is bicircular then it is succinctly represented by a description of a graph and a list of the balanced loops. 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. 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.

Let (U,V) be a partition of E such that $\lambda_M(U) \leqslant \lambda$ for some positive integer λ. Let N be the set of vertices that are in both $G[U]$ and $G[V]$, so that $|N| \leqslant 14\lambda - 12$ by Proposition 32. We will describe an equivalence relation $\approx_U$ on the subsets of U and then show that $\approx_U$ can be computed in time bounded by $O(\pi(\lambda)|E(M)|^c)$ for some function π and some constant c. Moreover, we will prove that $\approx_U$ satisfies conditions (i) and (ii) of Definition 14.

Let X be a subset of U. If M is H-gain-graphic then we can choose a maximal forest of $G[X]$ and then perform switching operations so that every edge in the forest is labelled with the identity of H [20, Lemma 5.3]. If M is bicircular, we choose the maximal forest but we do not need to perform any switchings. Now X is dependent in M if and only if there is an edge not in the spanning forest that receives an identity label, or if there are two distinct edges not in the maximal forest that are in the same component of $G[X]$. Thus we can test the independence of X in polynomial time.

Let X and X' be subsets of U. We consider the circumstances under which we declare

X and X' to be equivalent under $\approx_U$. Firstly, if X and X' are both dependent, then $X \approx_U X'$. If exactly one of X and X' is dependent, then $X \not\approx_U X'$. Now assume that both X and X' are independent. In this case, any component of $G[X]$ or $G[X']$ contains at most one cycle. Let u and v be vertices in a component of $G[X]$ or $G[X']$. Let Γ be this component. We claim there are at most two paths of Γ that join u to v. This is clear if Γ is a tree, so assume that Γ contains exactly one cycle. Let e be an edge of Γ such that $\Gamma\backslash e$ is a tree. If there are three distinct paths from u to v in Γ, then two of them use e. Now there must be two distinct paths of $\Gamma\backslash e$ from either u or v to an end-vertex of e. This is impossible so our claim is proved.

Let $\mathcal{X}$ and $\mathcal{X}'$ be the sets of connected components in $G[X]$ and $G[X']$, respectively, that have non-empty intersection with N. In the case that X and X' are both independent, we declare that $X \approx_U X'$ if there is a bijection $\theta\colon \mathcal{X} \to \mathcal{X}'$ such that the following statements hold for every $\Gamma \in \mathcal{X}$,

(i) $\Gamma \cap N = \theta(\Gamma) \cap N$,

(ii) Γ contains a cycle if and only if $\theta(\Gamma)$ contains a cycle, and

(iii) in the case that M is H-gain-graphic and u and v are vertices of $\Gamma \cap N$, there is a path of Γ from u to v with gain-value h if and only if there is a path of $\theta(\Gamma)$ from u to v with gain-value h.

It is clear that $\approx_U$ is an equivalence relation. Next we count the equivalence classes. Let $\tau(|N|)$ be the number of partitions of N where at most one block of the partition is allowed to be empty. We will think of the non-empty blocks in this partition as being the intersections of components of $G[X]$ with N. Thus $\tau(|N|)$ counts the number of possible such intersections. Note that τ is a non-decreasing function on the integers.

To choose an equivalence class of $\approx_U$, we first choose a non-empty collection of pairwise disjoint subsets of N. Since the size of N is at most $14\lambda - 12$, we can do this in at most $\tau(14\lambda - 12)$ ways. Next we choose whether each of these components has a cycle or not. The number of components that intersect N is at most $14\lambda - 12$, so we can make this choice in at most $2^{14\lambda-12}$ ways. Finally, for each pair $(u, v) \in N \times N$, we choose at most two gain-values in H for paths from u to v. The number of ways we can make this choice is at most $|H|^{2(14\lambda-12)^2}$. Thus the number of equivalence classes under $\approx$ is at most

$$\tau(14\lambda - 12)2^{14\lambda-12}|H|^{2(14\lambda-12)^2}.$$

Let this number be denoted by $\pi(\lambda)$. Since H is fixed, $\pi(\lambda)$ depends only on λ. It is clear that we can test the equivalence $X \approx_U X'$ in time bounded by $\pi(\lambda)|E(M)|^c$ for some constant c. Now we can complete the proof of Theorem 34 by showing 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 $Z \subseteq V$, and let C be a circuit of M contained in $X \cup Z$. We will prove that $X' \cup Z$ is also dependent in M and this will complete the proof.

First assume that C is a balanced cycle. If C is a balanced loop, then it is contained in Z, since X is independent. In this case $X' \cup Z$ is dependent and we have nothing left to prove. Therefore we assume that C is a balanced cycle with more than one edge, so M is an H-gain-graphic matroid. Now each component of $G[X \cap C]$ is a path, P, between two vertices of N. We can replace each such P with a path P' of $G[X']$ running between the same vertices. Moreover, since $X \approx X'$ holds, we can choose P' so that $\sigma(P') = \sigma(P)$. When we perform all these substitutions on C we obtain a walk W in $G[X' \cup Z]$, where $\sigma(W)$ is the identity of H. It is now easy to prove that $G[W]$ contains either a balanced cycle, or two distinct cycles. In either case $G[X' \cup Z]$ is dependent in M so we are done.

Now we assume that C is a theta subgraph or a handcuff. Let $\Gamma_1, \ldots, \Gamma_n$ be the connected components of $G[X]$ that contain edges of C. For each i let D_i be the set of edges of C contained in Γ_i. Note that $D_1, \ldots, D_n$ are pairwise disjoint sets of edges. Because $X \approx X'$ we can make a choice of D'_i, a minimal set of edges in X' for each i such that the following conditions hold.

(i) $G[D'_i]$ is connected,

(ii) every vertex of $G[D_i] \cap N$ is in $G[D'_i]$, and

(iii) $G[D'_i]$ contains a cycle if and only if $G[D_i]$ contains a cycle.

Let C_0 be C. For each i let C_i be the subgraph obtained from C_{i-1} by replacing the edges of D_i with D'_i. Thus C_n is a subgraph of $G[X' \cup Z]$. It is clear that each C_i is connected. We will show that C_n contains at least two cycles, and then we will be done.

For any graph, Γ, let $\nu(\Gamma)$ be $|E(\Gamma)| - |V(\Gamma)|$. If Γ is connected, then $\nu(\Gamma) \geqslant -1$. 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

$$\nu(\Gamma) = \nu(\Gamma[L]) + \nu(\Gamma[R]) + \gamma. \tag{3}$$

Note that $\nu(C_0) = 1$. We assume inductively that $\nu(C_{i-1}) \geqslant 1$. Note that $G[D_i]$ may not be connected, but $G[D'_i]$ is connected. Our choice of D'_i means that $\nu(D'_i) \geqslant \nu(D_i)$. Furthermore, $G[D'_i]$ has at least as many vertices in common with $G[C_{i-1} - D_i]$ as $G[D_i]$ does. It now follows from (3) that $\nu(C_i) \geqslant \nu(C_{i-1}) \geqslant 1$. Thus $\nu(C_n) \geqslant 1$, and since C_n is connected, it follows that C_n contains at least two cycles, as required. □

Corollary 35. *Let $\mathcal{M}$ be the class of bicircular or H-gain-graphic matroids (with H a finite group). Let ψ be any sentence in CMS_0. We can test whether matroids in $\mathcal{M}$ satisfy ψ using an algorithm that is fixed-parameter tractable with respect to branch-width.*

Proof. This will follow immediately from [8, Theorem 6.7] and Theorem 34 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 technical. 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 labelled 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 balanced loops after contracting e.

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. □

Corollary 35 completes the proof of Theorem 2.

Remark 36. Hliněný has shown [10, p. 348] that his work provides an alternative proof of Courcelle's Theorem. We can provide a simple new proof by relying on Corollary 35 and using bicircular matroids as models for graphs. We now briefly explain this strategy.

Let ψ be a sentence in the counting monadic second-order logic, CMS_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. Furthermore, we have predicates which allow us to assert that a set has cardinality p modulo q, for any appropriate choice of p and q. 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 $B(G^\circ)$. We let $\mathrm{Vert}(X_i)$ be the CMS_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 v is a vertex variable, we replace $\exists v$ with $\exists X_v\, \mathrm{Vert}(X_v) \wedge$, and we replace $\forall v$ with $\forall X_v\, \mathrm{Vert}(X_v)\ \rightarrow$. We perform a similar replacement for variables representing edges. If V is a variable representing a set of vertices, we replace $\exists V$ with

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

where $\mathsf{Sing}(X_1)$ is a predicate expressing that X_1 contains exactly one element. There are similar replacements for variables representing sets of edges and for universal quantifiers. Finally, we replace any occurrence of the predicate stating that e is incident with v with a CMS_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 35 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 subcubic 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 [5].

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 37. *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 $\{ab\colon a \in A, b \in B\}$ 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 $A \cup \{x\}$. □

Proposition 38. *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 $\mathrm{dw}(M) \leqslant K$ 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 $q - p \geqslant q/2 \geqslant p$ 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 37, we choose disjoint subsets $A = \{a_1, \ldots, a_q\}$ and $B = \{b_1, \ldots, b_q\}$ of H such that $a_i b_j \neq a_p b_q$ whenever $(i, j) \neq (p, q)$. Let AB be $\{a_i b_j \colon 1 \leqslant i, j \leqslant q\}$. We can 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 $\{a_i, b_j, a_i b_j\}$.

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 16, 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 $a_i b_j$ 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 $q/2 \geqslant p$ 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 $q - p \geqslant q/2$ 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 $(q^2 + 2q)/3 \leqslant |V_e|$, 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. □

9 Open problems

We have proved that the class of lattice path matroids is pigeonhole, but we have not yet proved that it is strongly pigeonhole. Nevertheless, we believe this to be the case.

Conjecture 39. The class of lattice path matroids is efficiently 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 22 (and Conjecture 39) exemplify a more general result.

Conjecture 40. The class of matroids that are both transversal and cotransversal is strongly pigeonhole.

Despite the existence of examples as in Remark 33, we firmly believe the next conjecture.

Conjecture 41. The class of bicircular matroids is efficiently pigeonhole. Let H be a finite group. The class of H-gain-graphic matroids is efficiently pigeonhole.

Acknowledgements

Daryl Funk and Dillon Mayhew are supported by a Rutherford Discovery Fellowship, managed by Royal Society Te Apārangi.

We thank Geoff Whittle for several important conversations, and Yves de Cornulier for suggesting a proof of Proposition 37. We also thank the referee for their helpful comments.

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**(1):#P106, (2011).

[4] J. E. Bonin and A. de Mier. Lattice path matroids: structural properties. *European J. Combin.* **27** (2006), no. 5, 701–738.

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

[6] B. Courcelle and J. Engelfriet. *Graph structure and monadic second-order logic*, volume 138 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, Cambridge (2012).

[7] R. G. Downey and M. R. Fellows. *Fundamentals of parameterized complexity.* Texts in Computer Science. Springer, London (2013).

[8] D. Funk, D. Mayhew, and M. Newman. Tree automata and pigeonhole classes of matroids: I. *Algorithmica* **84** (2022), no. 7, 1795–1834.

[9] 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).

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

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

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

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

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

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

[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.