TRANSACTIONS OF THE
AMERICAN MATHEMATICAL SOCIETY
Volume 370, Number 8, August 2018, Pages 5907–5929
http://dx.doi.org/10.1090/tran/7408
Article electronically published on April 17, 2018

YES, THE "MISSING AXIOM" OF MATROID THEORY IS LOST FOREVER

DILLON MAYHEW, MIKE NEWMAN, AND GEOFF WHITTLE

ABSTRACT. We prove there is no sentence in the monadic second-order language MS_0 that characterises when a matroid is representable over at least one field, and no sentence that characterises when a matroid is $\mathbb{K}$-representable, for any infinite field $\mathbb{K}$. By way of contrast, because Rota's Conjecture is true, there is a sentence that characterises $\mathbb{F}$-representable matroids, for any finite field $\mathbb{F}$.

1. INTRODUCTION

A matroid captures the notion of a discrete collection of points in space. Sometimes these points can be assigned coordinates in a consistent way, and sometimes they cannot. The problem of characterising when a matroid is *representable* has been the prime motivating force in matroid research since Whitney's founding paper [13].

Plenty of effort has been invested in characterising matroid representability via excluded minors. Less attention has been paid to the prospect of characterising representability via axioms. Perhaps this is because of Vámos's well-known article [12], which has been interpreted as stating that no such characterisation exists (see [4]). In [9], we pointed out that the possibility of characterising representable matroids in the language of Whitney's axioms was still open; that, in other words, we still did not know if "the missing axiom of matroid theory is lost forever", *contra* Vámos's title. We conjectured that in fact there was no such characterisation, and we made some partial progress towards resolving the conjecture by showing that it was impossible to characterise the class of representable matroids, or the class of matroids representable over an infinite field, using a logical language based on the rank function. However, that language imposed quite strong constraints on the form of quantification. In this article, we present a language with no such constraints, and we prove that it is impossible to characterise representability or representability over an infinite field in this more natural language. This is not to say that representability cannot be characterised in stronger languages: indeed, any language will suffice if it is strong enough to express the statement that the independent sets are in correspondence with the linearly independent sets of columns in a matrix.

The language that we develop here is a form of monadic second-order logic for matroids (similar to that used by Hliněný [6]), which we denote MS_0. As we show in Section 2, MS_0 is expressive enough to state the matroid axioms and to state when a matroid contains an isomorphic copy of a fixed minor. This means that any

Received by the editors January 11, 2015, and, in revised form, March 1, 2017.
2010 *Mathematics Subject Classification.* Primary 03C13, 05B35.

minor-closed class of matroids can be characterised with an MS_0 sentence as long as it has a finite number of excluded minors. In particular, since Rota's Conjecture has been positively resolved by Geelen, Gerards, and Whittle (see [5]), it follows that the class of $\mathbb{F}$-representable matroids can be characterised by a sentence in MS_0 whenever $\mathbb{F}$ is a finite field. Our main results show that this is not the case for infinite fields, nor is it possible to characterise the matroids that are representable over at least one field using an MS_0 sentence. When we say that a matroid is *representable* we mean it is representable over at least one field.

Theorem 1.1. *There is no sentence ψ in MS_0 such that a matroid is representable if and only if it satisfies ψ.*

Theorem 1.2. *Let $\mathbb{K}$ be any infinite field. There is no sentence $\psi_{\mathbb{K}}$ in MS_0 such that a matroid is $\mathbb{K}$-representable if and only if it satisfies $\psi_{\mathbb{K}}$.*

These theorems may seem stronger than those in [9], but in fact the results are independent of each other. The logical language used in [9] had constraints on quantification, unlike MS_0, but it also had access to the rank function and to the arithmetic of the integers, while MS_0 does not.

Theorems 1.1 and 1.2 follow easily from two lemmas. Let k be a positive integer. Let M_1 and M_2 be matroids. We will say that a *k-certificate* for M_1 and M_2 is a pair, (M', ψ), where M' is a matroid satisfying $E(M') \cap (E(M_1) \cup E(M_2)) = \emptyset$, and ψ is a sentence in MS_0 with k variables such that ψ is satisfied by exactly one of the direct sums $M_1 \oplus M'$ and $M_2 \oplus M'$. We define M_1 and M_2 to be *k-equivalent* if there is no k-certificate for M_1 and M_2. This relation is obviously reflexive and symmetric. Assume that M_1 is k-equivalent to M_2 and M_2 is k-equivalent to M_3, but that (M', ψ) is a k-certificate for M_1 and M_3. Relabelling the ground set of a matroid has no effect on whether it satisfies a sentence in MS_0. Therefore we can assume that $E(M')$ is disjoint from $E(M_1) \cup E(M_2) \cup E(M_3)$. Now (M', ψ) is a k-certificate for M_1 and M_2 or for M_2 and M_3, a contradiction. Therefore k-equivalence truly is an equivalence relation.

If two matroids are k-equivalent, then no k-variable sentence can distinguish them, even after adjoining an arbitrary matroid via a direct sum.

Lemma 1.3. *Let k be a positive integer. There are only finitely many equivalence classes of matroids under the relation of k-equivalence.*

In Section 3, we will find an explicit bound on the number of equivalence classes. By using Lemma 1.3, we can easily deduce Theorem 1.1.

Proof of Theorem 1.1. Assume that there is a sentence, ψ, in MS_0 that characterises representable matroids. Let k be the number of variables in ψ. We apply Lemma 1.3. Because there are infinitely many prime numbers, we can assume that M_1 and M_2 are k-equivalent, where $M_1 \cong \mathrm{PG}(2,p)$ and $M_2 \cong \mathrm{PG}(2,p')$ for distinct primes, p and p'. We choose M' to be isomorphic to M_1, where $E(M') \cap (E(M_1) \cup E(M_2)) = \emptyset$. Then ψ is satisfied by both $M_1 \oplus M'$ and $M_2 \oplus M'$, or it is satisfied by neither. But $M_1 \oplus M' \cong \mathrm{PG}(2,p) \oplus \mathrm{PG}(2,p)$ is representable over $\mathrm{GF}(p)$ [11, Proposition 4.2.11]. On the other hand, both $\mathrm{PG}(2,p')$ and $\mathrm{PG}(2,p)$ are isomorphic to minors of $M_2 \oplus M'$ [11, 4.2.19], so it follows from [11, Proposition 3.2.4] and [1, Proposition 7.3] that if $M_2 \oplus M'$ is representable over a field, then that field must simultaneously have subfields isomorphic to $\mathrm{GF}(p)$ and $\mathrm{GF}(p')$, an impossibility. To summarise, $M_1 \oplus M'$ is representable, and $M_2 \oplus M'$ is not, but ψ is satisfied by both or by neither. Thus ψ certainly does not characterise representable matroids. □

The notion of k-equivalence is reminiscent of the Myhill-Nerode characterisation of regular languages (see [10] or [3, Section 6.1]). Lemma 1.3 is also a matroid analogue of the fact that a graph property definable in monadic second-order logic can be recognised by an automaton [2] and is therefore *finite*, in the sense of Lengauer and Egon [7]. By way of contrast, the theorem in [9] used a proof technique that was essentially an Ehrenfeucht-Fraïssé game (see [3, Section 2.2]). Note that if two matroids are k-equivalent, then they satisfy exactly the same k-variable sentences (since the empty matroid is not a k-certificate). This implies the known fact that there are only finitely many *rank-k 0-types* (see [8, Section 3.4] for an explanation).

Our second lemma will be used to prove Theorem 1.2. In this case, it will not suffice to use direct sums, as the sum of two $\mathbb{K}$-representable matroids is also $\mathbb{K}$-representable. Thus we use the notion of a *proper amalgam* (which will be precisely defined in Section 4). Let $\mathcal{M}_\ell$ be the set of matroids that contain a $U_{2,5}$-restriction on the set $\ell = \{a, b, x, y, z\}$. If M_1 and M_2 are matroids in $\mathcal{M}_\ell$ and $E(M_1) \cap E(M_2) = \ell$, then the proper amalgam of M_1 and M_2 exists and is denoted by $\mathrm{Amal}(M_1, M_2)$. The ground set of $\mathrm{Amal}(M_1, M_2)$ is $E(M_1) \cup E(M_2)$, and $\mathrm{Amal}(M_1, M_2)|E(M_i) = M_i$, for $i = 1, 2$.

Let k be a positive integer. Let M_1 and M_2 be matroids in $\mathcal{M}_\ell$. A *(k,ℓ)-certificate* is a pair, (M', ψ), where $M' \in \mathcal{M}_\ell$ satisfies $E(M') \cap (E(M_1) \cup E(M_2)) = \ell$ and ψ is a k-variable sentence that is satisfied by exactly one of $\mathrm{Amal}(M_1, M')$ and $\mathrm{Amal}(M_2, M')$. We say that M_1 and M_2 are *(k,ℓ)-equivalent* if there is no such certificate.

Lemma 1.4. *Let k be a positive integer. There are only finitely many equivalence classes of $\mathcal{M}_\ell$ under the relation of (k,ℓ)-equivalence.*

Again, we will explicitly bound the number of equivalence classes.

In Section 5, we will construct two families of matroids in $\mathcal{M}_\ell$ by using *gain graphs*. Loosely speaking, a gain graph is a graph equipped with edge labels that come from a group. For each such graph, there is a corresponding gain-graphic matroid, whose ground set is the edge-set of the graph. Let $\mathbb{K}$ be a field, let $s, t \geq 3$ be integers, and let α and β be elements in $\mathbb{K} - \{0\}$ with orders greater than, respectively, s and $2t(t-1)$. For each such pair of tuples, $(\mathbb{K}, s, \alpha)$ and $(\mathbb{K}, t, \beta)$, there are unique gain graphs, which we will denote by $\Gamma(\mathbb{K}, \alpha, s)$ and $\Delta(\mathbb{K}, \beta, t)$. (We postpone the exact descriptions until Section 5.) The edge labels of $\Gamma(\mathbb{K}, \alpha, s)$ and $\Delta(\mathbb{K}, \beta, t)$ come from the multiplicative group of $\mathbb{K}$.

Assume that M corresponds to the gain graph $\Gamma(\mathbb{K}, \alpha, s)$ and that M' corresponds to $\Delta(\mathbb{K}, \beta, t)$. We also assume that $E(M) \cap E(M') = \ell$. In the case that $\alpha = \beta$, where the order of α is greater than $\max\{s, 2t(t-1)\}$, both M and M' can be represented over $\mathbb{K}$, but $\mathrm{Amal}(M, M')$ can be represented over $\mathbb{K}$ if and only if $s = t$. This means that Lemma 1.4 quickly leads to a proof of Theorem 1.2, with the two families of gain-graphic matroids playing the same role that projective planes did in the proof of Theorem 1.1. Details of the proof will be left until the end of the paper.

In fact, Lemma 1.4 is sufficient to prove both Theorems 1.1 and 1.2, since if $\mathrm{Amal}(M, M')$ is not representable over the field $\mathbb{K}$, then it is not representable over any field (Lemma 5.3). However, we feel that Lemma 1.3 is more intuitive, and also interesting in its own right, so we prefer to prove that lemma and then note the changes required to produce a proof of Lemma 1.4.

Lemma 1.4 also implies the following (unsurprising) facts: using MS_0 to characterise increasingly large finite fields requires increasingly large sentences. Furthermore, it is not possible to axiomatise the class of matroids representable over a given characteristic.

Corollary 1.5. *Let $\mathcal{Q}$ be the set of prime powers. For each $q \in \mathcal{Q}$, let ψ_q be an MS_0 sentence such that a matroid is* $\mathrm{GF}(q)$*-representable if and only if it satisfies ψ_q. There is no integer N such that every sentence in $\{\psi_q\}_{q\in\mathcal{Q}}$ has at most N variables.*

Corollary 1.6. *Let c be either* 0 *or a prime number. There is no sentence ψ_c in MS_0 such that a matroid is representable over a field of characteristic c if and only if it satisfies ψ_c.*

The paper is structured as follows: Section 2 introduces the MS_0 language for matroids and discusses its expressive power. Section 3 gives a proof of Lemma 1.3. In Section 4 we define the proper amalgam of matroids along a $U_{2,5}$-restriction and prove some of its properties. Section 5 introduces gain-graphic matroids and defines the two special classes of matroids. Finally, in Section 6, we prove Lemma 1.4 and complete the proof of Theorem 1.2 and Corollaries 1.5 and 1.6. For all matroid essentials we refer to Oxley [11].

2. Monadic second-order logic

In this section we give a formal definition of our monadic second-order language for matroids. The language MS_0 includes a countably infinite supply of variables, $X_1, X_2, X_3, \ldots$, along with the binary predicate, $\subseteq$; the unary predicates, Sing and Ind; as well as the standard connectives $\wedge$ and $\neg$; and the quantifier $\exists$.

We recursively define formulas in MS_0 and simultaneously define their sets of variables. The following statements define expressions known as *atomic formulas.*

1. $X_i \subseteq X_j$ is an atomic formula for any variables X_i and X_j, and $\mathrm{Var}(X_i \subseteq X_j) = \{X_i, X_j\}$.
2. $\mathrm{Sing}(X_i)$ is an atomic formula for any variable X_i, and $\mathrm{Var}(\mathrm{Sing}(X_i)) = \{X_i\}$.
3. $\mathrm{Ind}(X_i)$ is an atomic formula for any variable X_i, and $\mathrm{Var}(\mathrm{Ind}(X_i)) = \{X_i\}$.

A *formula* is an expression generated by a finite application of the following rules. Every formula has an associated set of variables and *free variables*:

1. Every atomic formula, ψ, is a formula, and $\mathrm{Fr}(\psi) = \mathrm{Var}(\psi)$.
2. If ψ is a formula, then $\neg\psi$ is a formula, and $\mathrm{Var}(\neg\psi) = \mathrm{Var}(\psi)$, while $\mathrm{Fr}(\neg\psi) = \mathrm{Fr}(\psi)$.
3. If ψ_1 and ψ_2 are formulas and $\mathrm{Fr}(\psi_i) \cap (\mathrm{Var}(\psi_j) - \mathrm{Fr}(\psi_j)) = \emptyset$ for $\{i, j\} = \{1, 2\}$, then $\psi_1 \wedge \psi_2$ is a formula, and $\mathrm{Var}(\psi_1 \wedge \psi_2) = \mathrm{Var}(\psi_1) \cup \mathrm{Var}(\psi_2)$, while $\mathrm{Fr}(\psi_1 \wedge \psi_2) = \mathrm{Fr}(\psi_1) \cup \mathrm{Fr}(\psi_2)$.
4. If ψ is a formula and $X_i \in \mathrm{Fr}(\psi)$, then $\exists X_i\psi$ is a formula, and $\mathrm{Var}(\exists X_i\psi) = \mathrm{Var}(\psi)$, while $\mathrm{Fr}(\exists X_i\psi) = \mathrm{Fr}(\psi) - \{X_i\}$.

A variable in $\mathrm{Var}(\psi)$ is *free* if it is in $\mathrm{Fr}(\psi)$, and *bound* otherwise. A formula is *quantifier-free* if all of its variables are free and is a *sentence* if all of its variables are bound. If ψ is a quantifier-free formula, then we will define the *depth* of ψ to be the number of applications of rules (2) and (3) required to construct ψ. Rule (3) insists that no variable can be free in one of ψ_1 and ψ_2 and bound in the other if $\psi_1 \wedge \psi_2$ is to be a formula. We can overcome this constraint if necessary by renaming the bound variables in a formula.

If ψ is a formula and $X_i \in \mathrm{Fr}(\psi)$, then we use $\forall X_i \psi$ as shorthand for $\neg(\exists X_i \neg\psi)$. We also use the shorthand $\psi_1 \vee \psi_2$ to mean $\neg((\neg\psi_1) \wedge (\neg\psi_2))$ and we use $\psi_1 \to \psi_2$ to mean $(\neg\psi_1) \vee \psi_2$. Likewise, we use $\psi_1 \leftrightarrow \psi_2$ to mean $(\psi_1 \to \psi_2) \wedge (\psi_2 \to \psi_1)$. We use $X \not\subseteq Y$ to stand for $\neg(X \subseteq Y)$.

Let ψ be a formula in MS_0. An *interpretation* of ψ is a pair (M, τ), where $M = (E, \mathcal{I})$ consists of a set, E, and a collection, $\mathcal{I}$, of subsets of E, and τ is a function from $\mathrm{Fr}(\psi)$ into the power set of E. We will recursively define what it means for (M, τ) to *satisfy* ψ, starting with the case that ψ is atomic. If ψ is $X_i \subseteq X_j$, then (M, τ) satisfies ψ if and only if $\tau(X_i) \subseteq \tau(X_j)$. If ψ is $\mathrm{Sing}(X_i)$, then (M, τ) satisfies ψ if and only if $|\tau(X_i)| = 1$. Finally, if ψ is $\mathrm{Ind}(X_i)$, then (M, τ) satisfies ψ if and only if $\tau(X_i)$ is in $\mathcal{I}$.

Now we assume that ψ is not atomic. If ψ is $\neg\phi$ for some formula ϕ, then (M, τ) satisfies ψ if and only if (M, τ) does not satisfy ϕ. Assume that ψ is $\phi_1 \wedge \phi_2$. Then (M, τ) satisfies ψ if and only if $(M, \tau\restriction_{\mathrm{Fr}(\phi_1)})$ satisfies ϕ_1 and $(M, \tau\restriction_{\mathrm{Fr}(\phi_2)})$ satisfies ϕ_2. Finally, assume that ψ is $\exists X_i \phi$, where X_i is a free variable in the formula ϕ. Then (M, τ) satisfies ψ if and only if there exists a subset, $Y_i \subseteq E$, such that the interpretation $(M, \tau \cup \{(X_i, Y_i)\})$ satisfies ϕ. If ψ is an MS_0 sentence, then we say that $M = (E, \mathcal{I})$ *satisfies* ψ (or ψ is satisfied by M) if the interpretation $(M, \emptyset)$ satisfies ψ.

We will spend some time illustrating the expressive power of MS_0. It is powerful enough to state the axioms for matroids and to characterise when a matroid contains a fixed minor.

If $t \geq 2$ is an integer, we use $\mathrm{Union}_t(X_{i_1}, \ldots, X_{i_t}, X_{i_{t+1}})$ as shorthand for the formula

$$\forall X\ \mathrm{Sing}(X) \to (X \subseteq X_{i_{t+1}} \leftrightarrow \bigvee_{1 \leq j \leq t} X \subseteq X_{i_j}).$$

The variable X stands for some variable different from each of $X_{i_1}, \ldots, X_{i_{t+1}}$. Clearly the formula $\mathrm{Union}_t(X_{i_1}, \ldots, X_{i_t}, X_{i_{t+1}})$ is satisfied by the interpretation (M, τ) if and only if $\tau(X_{i_{t+1}})$ is equal to $\tau(X_{i_1}) \cup \cdots \cup \tau(X_{i_t})$.

We let $\mathrm{Max}(X_i)$ stand for the formula

$$\mathrm{Ind}(X_i) \wedge (\forall X\ X_i \subseteq X \to (X \subseteq X_i \vee \neg\mathrm{Ind}(X))).$$

Therefore $\mathrm{Max}(X_i)$ is satisfied by τ in $M = (E, \mathcal{I})$ if and only if $\tau(X_i)$ is a maximal member of $\mathcal{I}$.

Let E be a finite set, and let $\mathcal{I}$ be a collection of subsets of E. Then $\mathcal{I}$ is the family of independent sets of a matroid, $M = (E, \mathcal{I})$, if and only if M satisfies the following sentences:

I1. $\exists X_1\ \mathrm{Ind}(X_1)$.
I2. $\forall X_1 \forall X_2\ (\mathrm{Ind}(X_1) \wedge (X_2 \subseteq X_1)) \to \mathrm{Ind}(X_2)$.
I3. $\forall X_1 \forall X_2\ (\mathrm{Max}(X_1) \wedge \mathrm{Ind}(X_2) \wedge \neg\mathrm{Max}(X_2))$
$\to \exists X_3\ \mathrm{Sing}(X_3) \wedge (X_3 \subseteq X_1) \wedge (X_3 \not\subseteq X_2)$
$\wedge \exists X_4\ (\mathrm{Union}_2(X_2, X_3, X_4) \wedge \mathrm{Ind}(X_4))$.

The sentence **I3** declares that if X_1 is a maximal set in $\mathcal{I}$ and X_2 is a non-maximal set, then there is an element $x \in X_1 - X_2$ such that $X_2 \cup \{x\}$ is in $\mathcal{I}$. It is not difficult to show that these axioms imply that the maximal members of $\mathcal{I}$ are equicardinal. From this it follows immediately that the maximal members of $\mathcal{I}$ obey the matroid basis axioms. Therefore **I1**, **I2**, and **I3** axiomatise matroids, as claimed.

Next we let N be a fixed matroid on the ground set $\{1,\ldots,n\}$, with $\mathcal{I}$ as its collection of independent sets. Let $\mathcal{D}$ be the set of dependent subsets of N. A matroid has a minor isomorphic to N if and only if it contains distinct elements $x_1,\ldots,x_n$ and an independent set, X_{n+1}, such that $\{x_1,\ldots,x_n\}\cap X_{n+1}=\emptyset$, and $\{x_{i_1},\ldots,x_{i_t}\}\cup X_{n+1}$ is independent precisely when $\{i_1,\ldots,i_t\}$ is an independent set of N. In this case, N is isomorphic to the minor produced by contracting X_{n+1} and restricting to the set $\{x_1,\ldots,x_n\}$. Thus we see that a matroid has a minor isomorphic to N if and only if it satisfies the following sentence:

$$\begin{aligned}\exists X_1\cdots\exists X_n\exists X_{n+1}\ &\mathrm{Ind}(X_{n+1})\wedge\bigwedge_{1\le i\le n}(\mathrm{Sing}(X_i)\wedge(X_i\not\subseteq X_{n+1}))\\ &\wedge\bigwedge_{1\le i<j\le n}X_i\not\subseteq X_j\\ &\wedge\bigwedge_{\{i_1,\ldots,i_t\}\in\mathcal{I}}(\exists X\ \mathrm{Union}_{t+1}(X_{i_1},\ldots,X_{i_t},X_{n+1},X)\wedge\mathrm{Ind}(X))\\ &\wedge\bigwedge_{\{i_1,\ldots,i_t\}\in\mathcal{D}}(\exists X\ \mathrm{Union}_{t+1}(X_{i_1},\ldots,X_{i_t},X_{n+1},X)\wedge\neg\mathrm{Ind}(X)).\end{aligned}$$

It follows that there is an MS_0 sentence that will characterise a minor-closed class of matroids, as long as that class has only finitely many excluded minors.

3. Proof of Lemma 1.3

Let k be a positive integer. Define $g_1(k,0)$ to be $2^{k(k+1)}3^k$, and recursively define $g_1(k,n+1)$ to be $2^{g_1(k,n)}$. Let $f_1(k)$ be $g_1(k,k)$. Our goal in this section is to prove Lemma 1.3. We restate the lemma here, with an explicit bound on the number of equivalence classes.

Lemma 3.1. *Let k be a positive integer. There are at most $f_1(k)$ equivalence classes of matroids under the relation of k-equivalence.*

Proof. We define a *registry* to be a $(k+2)\times k$ matrix with rows indexed by Ind, Sing, and $X_1,\ldots,X_k$, and columns indexed by $X_1,\ldots,X_k$. An entry in row Ind or in row X_i must be 'T' or 'F'. An entry in the row indexed by Sing is either '0', '1', or '>'. It follows that there are at most $g_1(k,0)$ possible registries.

We define a *depth*-0 *tree* to be a registry. Recursively, a *depth*-$(n+1)$ *tree* is a non-empty set of depth-n trees. An easy inductive argument shows that there are no more than $g_1(k,n+1)$ depth-$(n+1)$ trees, and hence no more than $f_1(k)$ depth-k trees.

A *stacked matroid* is a tuple $\mathcal{M}=(M,Y_1,\ldots,Y_m)$, where M is a matroid and each Y_i is a subset of $E(M)$. We define $||\mathcal{M}||$ to be m. We can identify the matroid M with the stacked matroid $\mathcal{M}=(M)$ and note that in this case, $||\mathcal{M}||=0$.

To each stacked matroid $\mathcal{M}$ satisfying $||\mathcal{M}||\le k$, we are going to associate a tree, $\mathcal{T}(\mathcal{M})$, of depth $k-||\mathcal{M}||$. We start by assuming that $k-||\mathcal{M}||=0$, so that $\mathcal{T}(\mathcal{M})$ is a depth-0 tree, which is to say, a registry. Let $\mathcal{M}$ be $(M,Y_1,\ldots,Y_k)$. For every j in $\{1,\ldots,k\}$, set the entry of $\mathcal{T}(\mathcal{M})$ in row Ind and column X_j to be 'T' if Y_j is independent in M, and otherwise set it to be 'F'. Now, for every pair

$i,j \in \{1,\ldots,k\}$, set the entry of $\mathcal{T}(\mathcal{M})$ in row X_i and column X_j to be 'T' if and only if $Y_i \subseteq Y_j$. Finally, for each $j \in \{1,\ldots,k\}$, set the entry of $\mathcal{T}(\mathcal{M})$ in row Sing and column X_j to be '0' if $|Y_j| = 0$, set it to be '1' if $|Y_i| = 1$, and set it to '$>$' otherwise. This defines $\mathcal{T}(\mathcal{M})$ in the case that $k - ||\mathcal{M}|| = 0$.

Now we make the inductive assumption that $\mathcal{T}(\mathcal{M})$ is defined when $k-||\mathcal{M}|| \leq n$, where n is some integer in $\{0,\ldots,k-1\}$. Let $\mathcal{M} = (M, Y_1,\ldots,Y_{k-n-1})$ be a stacked matroid. Thus $k - ||\mathcal{M}|| = n+1$. Let Y_{k-n} be any subset of $E(M)$. If $\mathcal{M}' = (M, Y_1,\ldots,Y_{k-n-1}, Y_{k-n})$, then $k - ||\mathcal{M}'|| = n$, so our inductive assumption means that $\mathcal{T}(\mathcal{M}')$ is defined and is a depth-n tree. Since a depth-$(n+1)$ tree is a non-empty set of depth-n trees, we simply define $\mathcal{T}(\mathcal{M})$ to be the set

$$\{\mathcal{T}(M, Y_1,\ldots,Y_{k-n-1}, Y_{k-n})\colon Y_{k-n} \subseteq E(M)\}.$$

We have now defined $\mathcal{T}(\mathcal{M})$ for each stacked matroid $\mathcal{M}$ that satisfies $||\mathcal{M}|| \leq k$. Note that if M is a matroid, then the stacked matroid $\mathcal{M} = (M)$ satisfies $||\mathcal{M}|| = 0$, and hence $\mathcal{T}(\mathcal{M})$ is a depth-k tree.

Let ψ be a formula in MS_0 such that either ψ is quantifier-free, or $\mathrm{Var}(\psi) = \{X_1,\ldots,X_k\}$. Let $b(\psi)$ be the number of bound variables in ψ. We are going to define what it means for $\mathcal{T}$ and $\mathcal{T}'$ to be *ψ-compatible* when $\mathcal{T}$ and $\mathcal{T}'$ are depth-$b(\psi)$ trees.

In the first case, we assume that ψ is quantifier-free, so that $b(\psi) = 0$, and $\mathcal{T}$ and $\mathcal{T}'$ are depth-0 trees, that is, registries. To start with, we assume that ψ is an atomic formula. If ψ is $X_i \subseteq X_j$, then we define $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible if and only if their entries in row X_i and column X_j are both 'T'. Similarly, if ψ is $\mathrm{Ind}(X_j)$, then we define $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible if and only if both $\mathcal{T}$ and $\mathcal{T}'$ have 'T' as their entries in row Ind and column X_j. Next we assume that ψ is $\mathrm{Sing}(X_j)$. Let ω be the entry of $\mathcal{T}$ in row Sing and column X_j. Let ω' be the analogous entry of $\mathcal{T}'$. We define $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible if and only if $\{\omega, \omega'\} = \{\text{'0'}, \text{'1'}\}$.

This defines ψ-compatibility in the case that ψ is atomic, so we will now assume it is not atomic. Since it is quantifier-free, this means that ψ has the form $\neg\phi$ or $\phi_1 \wedge \phi_2$. First assume that ψ is $\neg\phi$, where ϕ is quantifier-free. By induction on the depth of quantifier-free formulas, we can determine whether or not $\mathcal{T}$ and $\mathcal{T}'$ are ϕ-compatible. We define $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible if and only if $\mathcal{T}$ and $\mathcal{T}'$ are *not* ϕ-compatible. Next assume that ψ is $\phi_1 \wedge \phi_2$. Again, ϕ_1 and ϕ_2 have no bound variables, and by induction on the depth of quantifier-free formulas, we can determine whether $\mathcal{T}$ and $\mathcal{T}'$ are compatible relative to ϕ_1 and ϕ_2. We define $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible if and only if $\mathcal{T}$ and $\mathcal{T}'$ are both ϕ_1-compatible and ϕ_2-compatible. We have now defined ψ-compatibility in the case that ψ has no bound variables.

Next we will assume that $\mathrm{Var}(\psi) = \{X_1,\ldots,X_k\}$. By the previous paragraphs, we can make the inductive assumption that ψ-compatibility is defined if $b(\psi) \leq n$, where n is some integer in $\{0,\ldots,k-1\}$. Let ψ be a formula with $\mathrm{Var}(\psi) = \{X_1,\ldots,X_k\}$ and assume that ψ has $n+1$ bound variables. By renaming variables, we will assume that $\mathrm{Fr}(\psi) = \{X_1,\ldots,X_{k-n-1}\}$ and that $X_{k-n},\ldots,X_k$ are the bound variables of ψ. By standard techniques, we can assume that ψ is in *prenex normal form*. That is,

$$\psi = Q_{k-n}X_{k-n} \cdots Q_k X_k\ \psi',$$

where each Q_j is either $\exists$ or $\forall$, and ψ' is a quantifier-free formula in MS_0 with $\mathrm{Var}(\psi') = \{X_1, \ldots, X_k\}$. Let ϕ be the formula $Q_{k-n+1}X_{k-n+1} \cdots Q_k X_k\ \psi'$ obtained from ψ by removing the quantification of X_{k-n}.

Let $\mathcal{T}$ and $\mathcal{T}'$ be trees of depth $b(\psi) = n+1$. Thus $\mathcal{T}$ and $\mathcal{T}'$ are non-empty sets of depth-n trees. First consider the case that $Q_{k-n} = \exists$. The number of bound variables in ϕ is n. If $\mathcal{T}_0$ is a depth-n tree contained in $\mathcal{T}$ and $\mathcal{T}_0'$ is a depth-n tree in $\mathcal{T}'$, then by the inductive hypothesis, ϕ-compatibility of $\mathcal{T}_0$ and $\mathcal{T}_0'$ is defined. We define $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible if and only if there exist trees $\mathcal{T}_0 \in \mathcal{T}$ and $\mathcal{T}_0' \in \mathcal{T}'$ that are ϕ-compatible.

Similarly, if $Q_{k-n} = \forall$, we define $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible if and only if $\mathcal{T}_0$ and $\mathcal{T}_0'$ are ϕ-compatible for every tree $\mathcal{T}_0 \in \mathcal{T}$ and every tree $\mathcal{T}_0' \in \mathcal{T}'$. This completes the definition of ψ-compatibility.

The following claim contains the heart of the proof of Lemma 3.1.

Claim 3.1.1. Let ψ be an MS_0 formula such that either ψ is quantifier-free or $\mathrm{Var}(\psi) = \{X_1, \ldots, X_k\}$. If $\mathrm{Var}(\psi) = \{X_1, \ldots, X_k\}$, then let m be $|\mathrm{Fr}(\psi)|$ and assume that $\mathrm{Fr}(\psi) = \{X_1, \ldots, X_m\}$. Otherwise, let m be k. Let $\mathcal{M} = (M, Y_1, \ldots, Y_m)$ and $\mathcal{M}' = (M', Y_1', \ldots, Y_m')$ be stacked matroids, where $E(M) \cap E(M') = \emptyset$. Define τ to be the function that takes X_i to $Y_i \cup Y_i'$, for each $X_i \in \mathrm{Fr}(\psi)$. The interpretation $(M \oplus M', \tau)$ satisfies ψ if and only if the trees $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible.

Proof. Let $b(\psi)$ be the number of bound variables in ψ. We will prove the claim by induction on $b(\psi)$. Note that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ both have depth $k - m = b(\psi)$.

For our base case, we assume that $b(\psi) = 0$, so that ψ is quantifier-free, $||\mathcal{M}|| = ||\mathcal{M}'|| = k$, and both $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are registries. Start by assuming that ψ is an atomic formula. Consider the case that ψ is $X_i \subseteq X_j$. Then $(M \oplus M', \tau)$ satisfies ψ if and only if $\tau(X_i) \subseteq \tau(X_j)$, which is true if and only if $Y_i \subseteq Y_j$ and $Y_i' \subseteq Y_j'$. But this is the case if and only if $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ both contain 'T' in row X_i and column X_j. This is exactly what it means for $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ to be ψ-compatible, so we are done in this case.

In our next case, ψ is $\mathrm{Ind}(X_j)$. Then $(M \oplus M', \tau)$ satisfies ψ if and only if $\tau(X_j)$ is independent in $M \oplus M'$. This is true if and only if Y_j is independent in M and Y_j' is independent in M'. In turn, this is true if and only if $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ both contain 'T' in row Ind and column X_j, which is the case if and only if $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible.

Next, we assume that ψ is $\mathrm{Sing}(X_j)$. Then $(M \oplus M', \tau)$ satisfies ψ if and only if $|\tau(X_j)| = 1$, and this is true if and only if $\{|Y_j|, |Y_j'|\} = \{0, 1\}$. This holds if and only if the entries of $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ in row Sing and column X_j are '0' and '1', in some order. Once again, this is true precisely when $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible. We have finished the case that ψ is atomic, so now we assume that ψ is not atomic.

Since ψ is quantifier-free, it has the form $\neg\phi$ or $\phi_1 \wedge \phi_2$. Consider the former case. By induction on the depth of quantifier-free formulas, we can conclude that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ϕ-compatible if and only if $(M \oplus M', \tau)$ satisfies ϕ. The definition of compatibility means that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible if and only if they are not ϕ-compatible, which is the case exactly when $(M \oplus M', \tau)$ satisfies ψ.

Next we assume that ψ is $\phi_1 \wedge \phi_2$, where ϕ_1 and ϕ_2 have no bound variables. Again, we use induction on the depth of quantifier-free formulas. We conclude that $(M \oplus M', \tau\restriction_{\mathrm{Fr}(\phi_\alpha)})$ satisfies ϕ_α if and only if $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ϕ_α-compatible, for $\alpha = 1, 2$. This holds if and only if $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible. Thus we have proved the claim in the case that $b(\psi) = 0$.

We make the inductive assumption that the claim holds when the number of bound variables is at most n, for some integer $n \in \{0, \ldots, k-1\}$. Consider the case that $b(\psi) = n + 1$. We have assumed that $\mathrm{Fr}(\psi) = \{X_1, \ldots, X_{k-n-1}\}$, and we can also assume that

$$\psi = Q_{k-n}X_{k-n} \cdots Q_k X_k \ \psi',$$

where each Q_j is a quantifier and ψ' is quantifier-free and satisfies $\mathrm{Var}(\psi') = \{X_1, \ldots, X_k\}$. Let ϕ be $Q_{k-n+1}X_{k-n+1} \cdots Q_k X_k \ \psi'$.

Consider the case that $Q_{k-n} = \exists$. Then $(M \oplus M', \tau)$ satisfies ψ if and only if there are subsets $Y_{k-n} \subseteq E(M)$ and $Y'_{k-n} \subseteq E(M')$ such that

$$(M \oplus M', \tau \cup \{(X_{k-n}, Y_{k-n} \cup Y'_{k-n})\})$$

satisfies ϕ. By the inductive assumption, this holds if and only if

$$\mathcal{T}(M, Y_1, \ldots, Y_{k-n-1}, Y_{k-n})$$

and

$$\mathcal{T}(M', Y'_1, \ldots, Y'_{k-n-1}, Y'_{k-n})$$

are ϕ-compatible. Now $\mathcal{T}(M, Y_1, \ldots, Y_{k-n-1}, Y_{k-n})$ is a depth-n tree contained in the depth-$(n+1)$ tree $\mathcal{T}(\mathcal{M})$, and $\mathcal{T}(M', Y'_1, \ldots, Y'_{k-n-1}, Y'_{k-n})$ is similarly contained in $\mathcal{T}(\mathcal{M}')$. Thus the recursive definition of compatibility means that $(M \oplus M', \tau)$ satisfies ψ if and only if $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible, exactly as desired.

The case when $Q_{k-n} = \forall$ is similar. In this case, $(M \oplus M', \tau)$ satisfies ψ if and only if $(M \oplus M', \tau \cup \{(X_{k-n}, Y_{k-n} \cup Y'_{k-n})\})$ satisfies ϕ for every choice of subsets $Y_{k-n} \subseteq E(M)$ and $Y'_{k-n} \subseteq E(M')$. By induction, this is true if and only if $\mathcal{T}(M, Y_1, \ldots, Y_{k-n-1}, Y_{k-n})$ and $\mathcal{T}(M', Y'_1, \ldots, Y'_{k-n-1}, Y'_{k-n})$ are ϕ-compatible, for every choice of Y_{k-n} and Y'_{k-n}. This holds if and only if $\mathcal{T}_0$ and $\mathcal{T}'_0$ are ϕ-compatible, for all trees $\mathcal{T}_0 \in \mathcal{T}(\mathcal{M})$ and $\mathcal{T}'_0 \in \mathcal{T}(\mathcal{M}')$. This is exactly what it means for $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ to be ψ-compatible, so the proof is complete. □

Let M_1 and M_2 be two matroids, which we consider as stacked matroids $\mathcal{M}_1 = (M_1)$ and $\mathcal{M}_2 = (M_2)$. We complete the proof of Lemma 3.1 by showing that if the trees $\mathcal{T}(\mathcal{M}_1)$ and $\mathcal{T}(\mathcal{M}_2)$ are equal, then M_1 and M_2 are k-equivalent. This will imply that the number of equivalence classes is at most the number of depth-k trees, and we will be done. Thus we assume that $\mathcal{T}(\mathcal{M}_1) = \mathcal{T}(\mathcal{M}_2)$.

Let M' be any matroid with $E(M') \cap (E(M_1) \cup E(M_2)) = \emptyset$, and let $\mathcal{M}' = (M')$ be the corresponding stacked matroid. Let ψ be any MS_0 sentence with $\mathrm{Var}(\psi) = \{X_1, \ldots, X_k\}$. Then Claim 3.1.1 implies that $M_1 \oplus M'$ satisfies ψ if and only if $\mathcal{T}(\mathcal{M}')$ is ψ-compatible with $\mathcal{T}(\mathcal{M}_1) = \mathcal{T}(\mathcal{M}_2)$, which holds if and only if $M_2 \oplus M'$ satisfies ψ. Therefore no k-certificate exists for M_1 and M_2, so they are k-equivalent, exactly as desired. □

4. Amalgams

Let M_1 and M_2 be simple matroids with ground sets E_1 and E_2, rank functions r_1 and r_2, and closure operators cl_1 and cl_2. Let ℓ be $E_1 \cap E_2$, where we assume that $M_1|\ell = M_2|\ell$.A matroid, M, on the ground set $E_1 \cup E_2$ is an *amalgam* of M_1

and M_2 if $M|E_1 = M_1$ and $M|E_2 = M_2$. A matroid is *modular* if $r(F) + r(F') = r(F \cap F') + r(F \cup F')$ whenever F and F' are flats. If $M_1|\ell$ is a modular matroid, then [11, Theorem 11.4.10] implies that

$$r(X) = \min\{r_1(Y \cap E_1) + r_2(Y \cap E_2) - r_1(Y \cap \ell) \;:\; X \subseteq Y \subseteq E_1 \cup E_2\} \tag{1}$$

is the rank function of an amalgam of M_1 and M_2, known as the *proper amalgam*. We denote this amalgam by $\mathrm{Amal}(M_1, M_2)$. Every rank-2 matroid is modular. (To see this, note that either $r(F) = r(F') = 1$ or one of F and F' is contained in the other. Neither of these cases leads to a violation of modularity.) Henceforth, we consider only the case that $r_1(\ell) = 2$. This means that $M_1|\ell$ is modular, so that $\mathrm{Amal}(M_1, M_2)$ is defined.

Proposition 4.1. *Assume that M_i is a simple matroid with ground set E_i, rank function r_i, and closure operator cl_i, for $i = 1, 2$. Let $\ell = E_1 \cap E_2$, where $M_1|\ell = M_2|\ell$ and $r_1(\ell) = 2$. Let X be a subset of $E_1 \cup E_2$. If $X \cap E_1$ is dependent in M_1 or if $X \cap E_2$ is dependent in M_2, then X is dependent in $\mathrm{Amal}(M_1, M_2)$. If $X \cap E_1$ is independent in M_1 and $X \cap E_2$ is independent in M_2, then X is dependent in $\mathrm{Amal}(M_1, M_2)$ if and only if*

(i) *$\ell \subseteq \mathrm{cl}_1(X \cap E_1)$ and $r_2((X - E_1) \cup \ell) < r_2(X - E_1) + 2$,*
(ii) *$\ell \subseteq \mathrm{cl}_2(X \cap E_2)$ and $r_1((X - E_2) \cup \ell) < r_1(X - E_2) + 2$, or*
(iii) *there is an element $y \in \ell$ such that $y \in \mathrm{cl}_1(X - E_2) \cap \mathrm{cl}_2(X - E_1)$.*

Proof. If $X \cap E_1$ is dependent in M_1, then $X \cap E_1$ is dependent in $\mathrm{Amal}(M_1, M_2)$, since $\mathrm{Amal}(M_1, M_2)|E_1 = M_1$. By symmetry, X is dependent in $\mathrm{Amal}(M_1, M_2)$ if $X \cap E_1$ is dependent in M_1 or if $X \cap E_2$ is dependent in M_2. Henceforth we assume that $X \cap E_1$ is independent in M_1 and $X \cap E_2$ is independent in M_2.

Assume statement (i) holds. Let Y be $X \cup \ell$. Then

$$\begin{aligned}|X| &= |X \cap E_1| + |X - E_1| \\ &= r_1(X \cap E_1) + r_2(X - E_1) \\ &> r_1(X \cap E_1) + r_2((X - E_1) \cup \ell) - 2 \\ &= r_1(Y \cap E_1) + r_2(Y \cap E_2) - r_1(Y \cap \ell),\end{aligned}$$

so by (1), the rank of X in $\mathrm{Amal}(M_1, M_2)$ is less than $|X|$, as desired. By symmetric arguments, we see that if (i) or (ii) holds, then X is dependent in $\mathrm{Amal}(M_1, M_2)$.

Next we assume that (iii) holds. Since $X \cap E_1$ contains no circuits of M_1 it follows that y is not in X. If $X \cap \ell$ contains distinct elements, u and v, then by performing circuit elimination on $\{y, u, v\}$ and a circuit contained in $(X - E_2) \cup y$ that contains y, we obtain a circuit of M_1 contained in $X \cap E_1$. This contradiction means that $|X \cap \ell| \in \{0, 1\}$. Let Y be $X \cup y$. Then

$$\begin{aligned}|X| &= |X \cap E_1| + |X \cap E_2| - |X \cap \ell| \\ &= r_1(X \cap E_1) + r_2(X \cap E_2) - |X \cap \ell| \\ &= r_1(Y \cap E_1) + r_2(Y \cap E_2) - (r_1(Y \cap \ell) - 1) \\ &> r_1(Y \cap E_1) + r_2(Y \cap E_2) - r_1(Y \cap \ell).\end{aligned}$$

Again we see that X is dependent in $\mathrm{Amal}(M_1, M_2)$, and this completes the proof of the 'if' direction.

For the 'only if' direction, we assume that X is dependent in $\mathrm{Amal}(M_1, M_2)$. As $X \cap E_1$ is independent in M_1 and $X \cap E_2$ is independent in M_2, it follows that X is contained in neither E_1 nor E_2. There is some set Y such that $X \subseteq Y \subseteq E_1 \cup E_2$ and $|X| > r_1(Y \cap E_1) + r_2(Y \cap E_2) - r_1(Y \cap \ell)$. Assume that amongst all such sets, Y has been chosen so that it is as small as possible. If y is an element in $Y - (X \cup E_2)$, then we could replace Y with $Y - y$. Therefore no such element exists. By symmetry it follows that $Y - X \subseteq \ell$. If $Y = X$, then $Y \cap E_1$ is independent in M_1, and $Y \cap E_2$ is independent in M_2, so $|X| > r_1(Y \cap E_1) + r_2(Y \cap E_2) - r_1(Y \cap \ell) = |Y| = |X|$. This contradiction means that there is an element, y, in $Y - X$. The minimality of Y means that

$$r_1(Y \cap E_1) + r_2(Y \cap E_2) - r_1(Y \cap \ell) \\ < r_1((Y - y) \cap E_1) + r_2((Y - y) \cap E_2) - r_1((Y - y) \cap \ell).$$

It follows that y is in $\mathrm{cl}_1((Y - y) \cap E_1)$ and $\mathrm{cl}_2((Y - y) \cap E_2)$, but not $\mathrm{cl}_1((Y - y) \cap \ell)$. We combine the observations in this paragraph to deduce that $|X \cap \ell| < |Y \cap \ell| < 3$.

Assume that $|X \cap \ell| = 1$, so that $|Y \cap \ell| = 2$ and $Y = X \cup y$. Let x be the element in $X \cap \ell$. Since $\mathrm{cl}_1(X \cap E_1) = \mathrm{cl}_1((Y - y) \cap E_1)$ contains x and y, it contains ℓ. As y is in $\mathrm{cl}_2((X - E_1) \cup x)$, it follows that $r_2((X - E_1) \cup \ell) = r_2((X - E_1) \cup x) < r_2(X - E_1) + 2$. Therefore statement (i) holds.

Now we assume that $|X \cap \ell| = 0$. If $Y \cap \ell = \{y\}$, then $Y = X \cup y$, and y is in

$$\mathrm{cl}_1((Y - y) \cap E_1) \cap \mathrm{cl}_2((Y - y) \cap E_2) = \mathrm{cl}_1(X - E_2) \cap \mathrm{cl}_2(X - E_1),$$

so statement (iii) holds. Therefore we assume that $Y \cap \ell = \{y, y'\}$, and hence $Y = X \cup \{y, y'\}$. Earlier statements imply that

$$y \in \mathrm{cl}_1((X \cap E_1) \cup y') \cap \mathrm{cl}_2((X \cap E_2) \cup y') \quad \text{and} \\ y' \in \mathrm{cl}_1((X \cap E_1) \cup y) \cap \mathrm{cl}_2((X \cap E_2) \cup y).$$

If y is in neither $\mathrm{cl}_1(X \cap E_1)$ nor $\mathrm{cl}_2(X \cap E_2)$, then $r_1(Y \cap E_1) = r_1((X \cap E_1) \cup y) = r_1(X \cap E_1) + 1$, and similarly, $r_2(Y \cap E_2) = r_2(X \cap E_2) + 1$. But this means that

$$r_1(Y \cap E_1) + r_2(Y \cap E_2) - r_1(Y \cap \ell) = r_1(X \cap E_1) + r_2(X \cap E_2) = |X|,$$

which is a contradiction. Hence, by using symmetry, we can assume that y is in $\mathrm{cl}_1(X \cap E_1)$. This means that y', and hence ℓ, is contained in $\mathrm{cl}_1(X \cap E_1)$. Also,

$$r_2((X - E_1) \cup \ell) = r_2((X - E_1) \cup y) < r_2(X - E_1) + 2,$$

so statement (i) holds, and the proof is complete. □

5. Gain-graphic matroids

In this section we introduce two families of matroids via gain graphs. Let G be an undirected graph (possibly containing loops and multiple edges) with edge-set E and vertex set V. Define $A(G)$ to be the following subset of $E \times V \times V$:

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

A *gain graph* (over the group H) is a pair (G, σ), where G is a graph and σ is a function from $A(G)$ to H, such that $\sigma(e, u, v) = \sigma(e, v, u)^{-1}$ for every non-loop edge e with end-vertices u and v. We say that σ is a *gain function*. If $C = v_0 e_0 v_1 e_2 \cdots e_t v_{t+1}$ is a cycle of G, where $v_0 = v_{t+1}$, then $\sigma(C)$ is defined to be

$$\sigma(e_0, v_0, v_1)\sigma(e_1, v_1, v_2) \cdots \sigma(e_t, v_t, v_{t+1}).$$

Note that, in general, H may be non-abelian, and the value of $\sigma(C)$ depends on the choice of starting point and orientation for C; however, if $\sigma(C)$ is equal to the identity of H, then this equality will hold no matter which starting point and orientation we choose. In this case, we say that C is *balanced*. A cycle that is not balanced is *unbalanced*.

The *gain-graphic* matroid $M(G,\sigma)$ has the edge set of G as its ground set. The circuits of $M(G,\sigma)$ are exactly the edge sets of balanced cycles, along with the minimal edge-sets that induce connected subgraphs containing at least two unbalanced cycles and no balanced cycles. Any such subgraph is either a *theta graph*, a *loose handcuff*, or a *tight handcuff*. A theta graph consists of two vertices joined by three internally disjoint paths, a loose handcuff consists of two vertex-disjoint cycles joined by a single path that intersects the cycles only in its end-vertices, and a tight handcuff consists of two edge-disjoint cycles that share exactly one vertex.

Assume that (G,σ) is a gain graph, where σ takes $A(G)$ to the multiplicative group of a field $\mathbb{K}$. Let $v_1,\ldots,v_m$ and $e_1,\ldots,e_n$ be orderings of the vertex and edge sets of G. We define a matrix, $D(G,\sigma)$, with entries from $\mathbb{K}$. The columns of $D(G,\sigma)$ are labelled by $e_1,\ldots,e_n$. Let $b_1,\ldots,b_m$ be the standard basis vectors. Assume that e_i is incident with vertices v_j and v_k, where $j \leq k$. (If e_i is a loop, then $j = k$.) The column labelled by e_i is equal to $b_j - \sigma(e_i,v_j,v_k)b_k$. Note that if e_i is a balanced loop, then column e_i is the zero vector, and if e_i is an unbalanced loop, then the column contains a single non-zero entry.

Proposition 5.1 (Theorem 2.1 of [14])**.** *Let (G,σ) be a gain graph over the multiplicative group of the field $\mathbb{K}$. The matrix $D(G,\sigma)$ represents the matroid $M(G,\sigma)$ over $\mathbb{K}$.*

Next we construct two families of gain graphs. Let $\mathbb{K}$ be a field. The gain functions of the two families will be into the multiplicative group of $\mathbb{K}$. Let $s \geq 3$ be an integer, and let α be an element in $\mathbb{K} - \{0\}$ with order greater than s. The gain graph $\Gamma(\mathbb{K},s,\alpha)$ has vertex set $\{u_1,\ldots,u_{s+1}\}$. Each vertex u_i in $\{u_2,\ldots,u_s\}$ is incident with a loop, a_i. In addition, u_1 is incident with the loop a, and u_{s+1} is incident with the loop b. The parallel edges x_i and y_i join u_i and u_{i+1} for each i in $\{1,\ldots,s\}$. Moreover, the edges x, y, and z join u_1 and u_{s+1}. We define the gain function σ so that it takes each loop to α and each x_i to 1. Furthermore, $\sigma(y_i,u_i,u_{i+1}) = \alpha$ for each i in $\{1,\ldots,s\}$, while $\sigma(x,u_1,u_{s+1}) = 1$, $\sigma(y,u_1,u_{s+1}) = \alpha^{s-1}$, and $\sigma(z,u_1,u_{s+1}) = \alpha^s$.

Now let $t \geq 3$ be an integer. We let β be an element in $\mathbb{K} - \{0\}$ with order greater than $2t(t-1)$. We construct the gain graph $\Delta(\mathbb{K},t,\beta)$. It has $\{v_1,\ldots,v_{2t}\}$ as its vertex set. Each vertex $v_i \in \{v_2,\ldots,v_{2t-1}\}$ is incident with a loop, b_i, while v_1 is incident with the loop a, and v_{2t} is incident with the loop b. For each $i \in \{1,\ldots,2t-1\}$, the edges e_i and f_i join v_i to v_{i+1}. The edges x, y, z, and g join the vertices v_1 and v_{2t}. The gain function σ takes each loop to β and each edge e_i to 1. The triple (f_i,v_i,v_{i+1}) is taken to β^{t-1} when $i \in \{1,\ldots,t\}$, and to β^t when $i \in \{t+1,\ldots,2t-1\}$. Thus t of the edges $f_1,\ldots,f_{2t-1}$ receive the label β^{t-1}, and the other $t-1$ receive the label β^t. The values of $\sigma(x,v_1,v_{2t})$, $\sigma(y,v_1,v_{2t})$, $\sigma(z,v_1,v_{2t})$, and $\sigma(g,v_1,v_{2t})$ are 1, β^{t-1}, β^t, and $\beta^{t(t-1)}$, respectively.

Figure 1 shows $\Gamma(\mathbb{K},s,\alpha)$ and $\Delta(\mathbb{K},t,\beta)$. The edge labels of loops have been omitted. Every edge label corresponds to the orientation of the edge shown in the drawing.

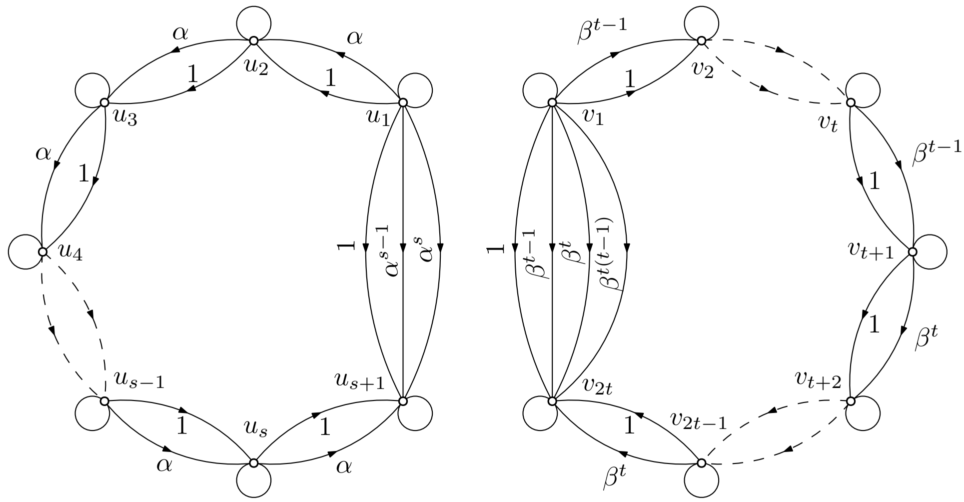

FIGURE 1. The gain graphs $\Gamma(\mathbb{K}, s, \alpha)$ and $\Delta(\mathbb{K}, t, \beta)$.

Note that whenever $M = M(\Gamma(\mathbb{K}, s, \alpha))$ and $M' = M(\Delta(\mathbb{K}, t, \beta))$, then $E(M) \cap E(M') = \{a, b, x, y, z\}$. Let ℓ be this intersection. Then $M|\ell$ and $M'|\ell$ are both isomorphic to $U_{2,5}$, so the discussion in Section 4 implies that $\mathrm{Amal}(M, M')$ is defined.

If G is a graph and X is a set of edges, then $G[X]$ denotes the subgraph of G containing the edges in X and all vertices that are incident with at least one edge in X.

Lemma 5.2. *Let $\mathbb{K}$ be a field, let $s \geq 3$ be an integer, and let α be an element in $\mathbb{K} - \{0\}$ with order greater than $2s(s-1)$. Let M be $M(\Gamma(\mathbb{K}, s, \alpha))$ and let M' be $M(\Delta(\mathbb{K}, s, \alpha))$. Then $\mathrm{Amal}(M, M')$ is $\mathbb{K}$-representable.*

Proof. Let ℓ be $\{a, b, x, y, z\}$. Let (G, σ) stand for the signed graph $\Gamma(\mathbb{K}, s, \alpha)$, and let (G', σ') stand for $\Delta(\mathbb{K}, s, \alpha)$. The lemma will follow from Proposition 5.1 if we can prove that $\mathrm{Amal}(M, M')$ is gain-graphic over the multiplicative group of $\mathbb{K}$. To this end, we construct a graph, H, by gluing together G and G'. We identify the vertices u_1 and v_1 as the new vertex w and identify u_{s+1} and v_{2s} as w'. The edge-set of H is exactly the union of the edge-sets of G and G'. Any edge incident with u_1 or v_1 in G or G' is incident with w in H, and any edge incident with u_{s+1} or v_{2s} is incident with w' in H. All other incidences are exactly as in G or G'. Let e be an edge of G or G', and let u and v be the vertices incident with e. (It may be the case that $u = v$.) If u is u_1 or v_1, then let $\hat{u}$ be w, and if $u = u_{s+1}$ or v_{2s}, then let $\hat{u}$ be w'. Otherwise define $\hat{u}$ to be u. We define $\hat{v}$ in exactly the same way. We define the function θ so that $\theta(e, \hat{u}, \hat{v}) = \sigma(e, u, v)$ if e is an edge of G and $\theta(e, \hat{u}, \hat{v}) = \sigma'(e, u, v)$ if e is an edge of G'. It is clear that θ is a well-defined gain function for H.

Let N be the gain-graphic matroid $M(H, \theta)$. We can prove the lemma by checking that N and $\mathrm{Amal}(M, M')$ are equal. We do this by showing that a set, X, is dependent in N if and only if it is dependent in $\mathrm{Amal}(M, M')$. Note that N is obviously an amalgam of M and M'.

For the first direction, we assume that X is a circuit in N. As N is an amalgam of M and M', we assume that X is contained in neither $E(M)$ nor $E(M')$. We start by considering the case that X is a balanced cycle in (H, θ). If X contains an edge joining w and w', then this edge is g, and $H[X]$ contains a path with vertex sequence $w, u_2, u_3, \ldots, u_s, w'$, for otherwise X is contained in $E(M)$ or $E(M')$. The product of edge labels along this path is α^j, where $j \leq s$. We also require that $\alpha^j = \alpha^{s(s-1)}$, since g is labelled with $\alpha^{s(s-1)}$, and X is a balanced cycle. But $\alpha^j = \alpha^{s(s-1)}$ cannot hold, as α has order greater than $2s(s-1)$, and $s \geq 3$ so $s(s-1) > s$. Therefore we conclude that X does not contain any edge between w and w', and since X is not contained in $E(M)$ or $E(M')$, it follows that it is the edge-set of a Hamiltonian cycle. Let α^j be the product of edge labels along the path in $H[X]$ with vertex sequence $w, u_2, u_3, \ldots, u_s, w'$. Thus $0 \leq j \leq s$. Let $\alpha^{p(s-1)+qs}$ be the product of edge labels along the path in $H[X]$ with vertex sequence $w, v_2, v_3, \ldots, v_{2s-1}, w'$, where p and q are non-negative integers satisfying $0 \leq p \leq s$ and $0 \leq q \leq s-1$. Thus $0 \leq p(s-1) + qs \leq 2s(s-1)$ and $\alpha^j = \alpha^{p(s-1)+qs}$, as X is balanced. As the order of α is greater than $2s(s-1)$, we deduce that $j = p(s-1) + qs$, and hence j is equal to 0, $s-1$, or s. In these three cases, x, y, or z is an element in $\mathrm{cl}_M(X - E(M')) \cap \mathrm{cl}_{M'}(X - E(M))$. Thus statement (iii) of Proposition 4.1 holds, so X is dependent in $\mathrm{Amal}(M, M')$.

Now we can assume that X does not contain a balanced cycle of (H, θ). Thus $H[X]$ is a theta graph or a handcuff. Let $\{M_1, M_2\}$ be $\{M, M'\}$. Assume that $H[X - E(M_1)]$ is a path from w to w'. None of the internal vertices of this path have degree three or more in $H[X]$. It follows that, regardless of whether $H[X]$ is a theta graph or a handcuff, $H[X \cap E(M_1)]$ contains an unbalanced cycle joined by a path to the loop a and an unbalanced cycle joined by a path to the loop b. Therefore $\{a, b\}$ (and hence all of ℓ) is contained in $\mathrm{cl}_{M_1}(X \cap E(M_1))$. Also, $H[(X - E(M_1)) \cup \{a, b\}]$ is a handcuff, and hence $(X - E(M_1)) \cup \{a, b\}$ is a circuit of M_2 that spans ℓ. This means that $r_{M_2}((X - E(M_1)) \cup \ell) < r_{M_2}(X - E(M_1)) + 2$. Now (i) of Proposition 4.1 holds, so X is dependent in $\mathrm{Amal}(M, M')$.

We can now assume that neither $H[X - E(M)]$ nor $H[X - E(M')]$ is a path from w to w'. Assume $H[X - E(M)]$ is a forest. As $H[X]$ has no vertices of degree one, the forest must be a path, and its end-vertices must be w and w', contradicting our assumption. By symmetry, it follows that each of $H[X - E(M)]$ and $H[X - E(M')]$ contains an unbalanced cycle. Since $H[X]$ is connected, either w or w' is on a path from one of these cycles to the other. Let us assume the former, since the latter case is identical. Now a is in a handcuff in $G[(X - E(M')) \cup a]$ and hence in a circuit of M that is contained in $(X - E(M')) \cup a$. By symmetry, a is also contained in a circuit of M' that is contained in $(X - E(M)) \cup a$. Thus statement (iii) of Proposition 4.1 holds, and X is dependent in $\mathrm{Amal}(M, M')$. We have proved that if X is dependent in N, it is dependent in $\mathrm{Amal}(M, M')$.

For the other direction, we assume that X is independent in N. This means that $H[X]$ contains no balanced cycles, and any connected component of $H[X]$ contains at most one cycle. Let us assume for a contradiction that X is dependent in $\mathrm{Amal}(M, M')$. In fact, we can assume that X is a circuit of $\mathrm{Amal}(M, M')$. As

N is an amalgam of M and M', it follows that neither $X \cap E(M)$ nor $X \cap E(M')$ is dependent, so X is contained in neither $E(M)$ nor $E(M')$. One of the three statements in Proposition 4.1 must hold.

We prove the following statements for M and M' simultaneously by letting $\{M_1, M_2\}$ be $\{M, M'\}$.

Claim 5.2.1. The subgraph $H[X - E(M_1)]$ contains either a connected component that contains both w and w', or a connected component that contains a cycle and at least one of w and w'.

Proof. Assume the claim is false, so that any component of $H[X - E(M_1)]$ contains at most one of w and w', and any component containing one of these vertices contains no cycle. This means that if p and q are distinct elements of ℓ, then $H[(X - E(M_1)) \cup \{p, q\}]$ contains no balanced cycles and no theta graphs or handcuffs. From this it follows that $\mathrm{cl}_{M_2}(X - E(M_1))$ does not contain any element of ℓ, so statement (iii) of Proposition 4.1 does not hold. Moreover, $r_{M_2}((X - E(M_1)) \cup \ell) = r_{M_2}(X - E(M_1)) + 2$. As one of the three statements in Proposition 4.1 must hold, it follows that ℓ is in $\mathrm{cl}_{M_2}(X \cap E(M_2))$ and $r_{M_1}((X - E(M_2)) \cup \ell) < r_{M_1}(X - E(M_2)) + 2$. But this now means that X contains at least two elements of ℓ, or else $\ell \nsubseteq \mathrm{cl}_{M_2}(X \cap E(M_2))$. Hence $\ell \subseteq \mathrm{cl}_{M_2}(X \cap \ell)$, so Proposition 4.1 implies that $X \cap E(M_1)$ is dependent in $\mathrm{Amal}(M, M')$. Since we have assumed X is a circuit of $\mathrm{Amal}(M, M')$, this means that $X \subseteq E(M_1)$, contrary to the hypothesis. Therefore Claim 5.2.1 holds. □

Claim 5.2.2. There is no connected component of $H[X - E(M_1)]$ that contains both w and w'.

Proof. Assume that H_0 is such a component. Then H_0 is contained in a connected component, H_1, of $H[X \cap E(M_2)]$. If H_1 contains a cycle, then by applying Claim 5.2.1 to $H[X - E(M_2)]$, we can deduce that the union of H_1 with a component of $H[X - E(M_2)]$ contains a theta graph or a handcuff. We have assumed that $H[X]$ contains no such subgraph, so this is a contradiction. Therefore H_1 contains no cycle, from which we deduce that $H[X - E(M_1)]$ is a path from w to w' and $X \cap \ell = \emptyset$. Note that $H[(X - E(M_1)) \cup a]$ contains no circuit of M_2, so $\ell \nsubseteq \mathrm{cl}_{M_2}(X - E(M_1))$.

If there is a component of $H[X - E(M_2)]$ that contains w and w', then by the reasoning in the previous paragraph, $H[X - E(M_2)]$ is a path from w to w', and $\ell \nsubseteq \mathrm{cl}_{M_1}(X - E(M_2))$. Therefore $H[X]$ is a Hamiltonian cycle, and the only statement in Proposition 4.1 that can hold is statement (iii). We have noted that a (and by symmetry b) is not in $\mathrm{cl}_{M_2}(X - E(M_1))$, so there is an edge, p, joining w and w' such that p is in both $\mathrm{cl}_{M_2}(X - E(M_1))$ and $\mathrm{cl}_{M_1}(X - E(M_2))$. This means that $H[(X - E(M_1)) \cup p]$ and $H[(X - E(M_2)) \cup p]$ are both balanced cycles. Thus the product of edge labels on the path $H[X - E(M_2)]$ is the inverse of the product on the path $H[X - E(M_1)]$ (assuming that we travel in a consistent direction around the Hamiltonian cycle $H[X]$). Hence X is a balanced cycle, a contradiction. Therefore no component of $H[X - E(M_2)]$ contains w and w'.

Recall that $X \cap \ell = \emptyset$ and neither a nor b is in $\mathrm{cl}_{M_2}(X - E(M_1))$. As one of the statements from Proposition 4.1 must hold, either there is an edge between w and w' that is in both $\mathrm{cl}_{M_2}(X - E(M_1))$ and $\mathrm{cl}_{M_1}(X - E(M_2))$ or $\mathrm{cl}_{M_1}(X - E(M_2))$ contains ℓ. Therefore in either case we can let p be an edge between w and w' that is in $\mathrm{cl}_{M_1}(X - E(M_2))$. Let C be a circuit of M_1 contained in $(X - E(M_2)) \cup p$

that contains p. No component of $H[X - E(M_2)]$ contains w and w' so $H[C-p]$ is not connected. It follows that $H[C]$ is a loose handcuff and p is an edge in the path between the two cycles. Therefore $H[X-E(M_2)]$ contains two distinct components, each containing a cycle and one of w and w'. As $H[X - E(M_1)]$ is a path from w to w', it follows that $H[X]$ is a handcuff, a contradiction. $\square$

We have shown that neither $H[X - E(M)]$ nor $H[X - E(M')]$ contains a component that contains w and w'. By using Claim 5.2.1, symmetry, and the fact that X contains no handcuffs, we can assume the following: there is a component of $H[X - E(M_2)]$ that contains w and a cycle, and any component that contains w' contains no cycle; similarly, there is a component of $H[X-E(M_1)]$ that contains w' and a cycle, and any component that contains w contains no cycle. It follows from this assumption (and the fact that $H[X]$ contains no handcuffs) that $X \cap \ell = \emptyset$. Notice that a is the only element in

$$\mathrm{cl}_{M_1}(X - E(M_2)) \cap \ell = \mathrm{cl}_{M_1}(X \cap E(M_1)) \cap \ell.$$

Similarly, $\mathrm{cl}_{M_2}(X \cap E(M_2)) \cap \ell = \{b\}$. Therefore none of the statements in Proposition 4.1 can hold, so we have a contradiction.

Now it follows that if X is independent in N, then it is also independent in $\mathrm{Amal}(M, M')$, so $N = \mathrm{Amal}(M, M')$, exactly as desired. This completes the proof of Lemma 5.2. $\square$

Lemma 5.3. *Let $\mathbb{K}$ be a field and let s and t be distinct integers satisfying $s, t \geq 3$. Let α be an element in $\mathbb{K} - \{0\}$ with order greater than $\max\{s, 2t(t-1)\}$. Let M be $M(\Gamma(\mathbb{K}, s, \alpha))$ and let M' be $M(\Delta(\mathbb{K}, t, \alpha))$. Then $\mathrm{Amal}(M, M')$ is not representable over any field.*

Proof. Let us assume that the matrix D represents $\mathrm{Amal}(M, M')$ over the field $\mathbb{L}$. Let B be the set $\{a_2, \ldots, a_s, a, b, b_2, \ldots, b_{2t-1}\}$. Thus B is the set of all loops in $\Gamma(\mathbb{K}, s, \alpha)$ and $\Delta(\mathbb{K}, t, \alpha)$. It is clear that $B \cap E(M)$ and $B \cap E(M')$ are independent in M and M', and moreover, $r_M((B - E(M')) \cup \ell) = r_M(B - E(M')) + 2$ and $r_{M'}((B - E(M)) \cup \ell) = r_{M'}(B - E(M)) + 2$. Now it follows easily from Proposition 4.1 that B cannot be dependent in $\mathrm{Amal}(M, M')$. If e is any element of the ground set of $\mathrm{Amal}(M, M')$ that is not in B, then $B \cup e$ contains a circuit of either M or M', and this circuit has cardinality three. From this it follows that B is a basis of $\mathrm{Amal}(M, M')$. We can assume that the columns of D labelled by the elements of B form an identity matrix. As the fundamental circuits relative to B all have cardinality three, every column of D contains either one or two non-zero elements. By scaling, we can assume that the first non-zero entry in each column is 1. Thus $D = D(G, \theta)$, for some gain graph (G, θ) over the multiplicative group of $\mathbb{L}$. By examining the fundamental circuits relative to B, we see that G is the graph in Figure 2.

By scaling rows of D, we can assume that

$$\theta(x_1, w, u_2) = \theta(x_s, u_s, w') = \theta(e_i, w, v_2) = 1.$$

Moreover, we can also assume that $\theta(x_i, u_i, u_{i+1}) = 1$ for each $i = 2, \ldots, s-1$ and that $\theta(e_i, v_i, v_{i+1}) = 1$ for each $i = 2, \ldots, 2t-2$. Note that $\{x_1, \ldots, x_s, x\}$ is a balanced cycle in $\Gamma(\mathbb{K}, s, \alpha)$ and that $\{e_1, \ldots, e_{2t-1}, x\}$ is a balanced cycle in $\Delta(\mathbb{K}, t, \alpha)$. It now follows from Proposition 4.1 that $\{x_1, \ldots, x_s, e_1, \ldots, e_{2t-1}\}$ is dependent in $\mathrm{Amal}(M, M')$, and we deduce that it is the edge-set of a balanced cycle in (G, θ). This in turn implies that $\theta(e_{2t-1}, v_{2t-1}, w') = 1$.

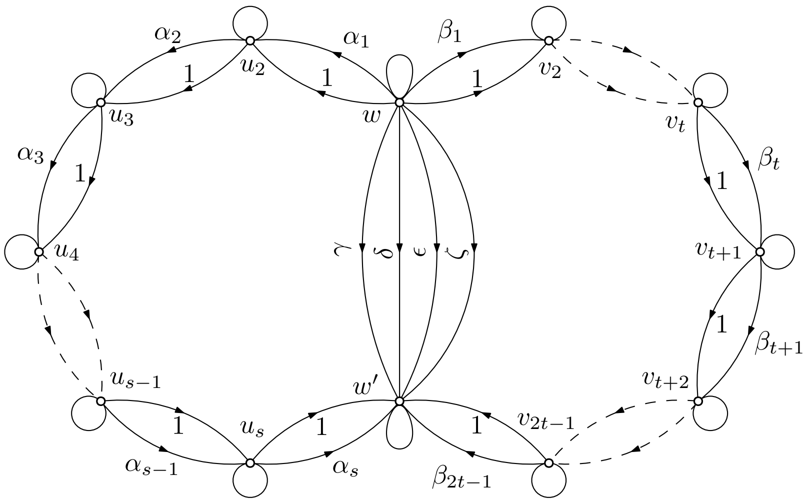

FIGURE 2. The gain graph (G, θ).

For $i \in \{2, \ldots, s-1\}$, let α_i be the value $\theta(y_i, u_i, u_{i+1})$. Define α_1 to be $\theta(y_1, w, u_2)$ and α_s to be $\theta(y_s, u_s, w')$. Similarly, for $i \in \{2, \ldots, 2t-2\}$, let β_i be $\theta(f_i, v_i, v_{i+1})$. Define β_1 to be $\theta(f_1, w, v_2)$, and let β_{2t-1} be $\theta(f_{2t-1}, v_{2t-1}, w')$. Let γ, δ, ϵ, and ζ be $\theta(x, w, w')$, $\theta(y, w, w')$, $\theta(z, w, w')$, and $\theta(g, w, w')$, respectively.

Because $\{x_1, \ldots, x_s, x\}$ is a balanced cycle in $\Gamma(\mathbb{K}, s, \alpha)$ and hence a circuit in $\mathrm{Amal}(M, M')$, it follows that it is also a balanced cycle in (G, θ). This means that $\gamma = 1$. Next we notice that $(\{y_1, \ldots, y_s\} - y_i) \cup \{x_i, y\}$ is a balanced cycle of $\Gamma(\mathbb{K}, s, \alpha)$ and hence of (G, θ), for any i in $\{1, \ldots, s\}$. The product of edge labels on this cycle in (G, θ) is $\alpha_1 \cdots \alpha_s \alpha_i^{-1} \delta^{-1}$, which implies that $\alpha_i = \alpha_1 \cdots \alpha_s \delta^{-1}$ for any $i \in \{1, \ldots, s\}$. Let α stand for $\alpha_1 \cdots \alpha_s \delta^{-1}$, so that $\alpha_i = \alpha$ for any $i \in \{1, \ldots, s\}$ and $\delta = \alpha^{s-1}$. As $\{y_1, \ldots, y_s, z\}$ is a balanced cycle, it follows that $\epsilon = \alpha^s$.

Next we observe that $(\{e_1, \ldots, e_{2t-1}\} - e_i) \cup \{f_i, y\}$ is a balanced cycle in $\Delta(\mathbb{K}, t, \alpha)$, and hence in (G, θ), for any $i \in \{1, \ldots, t\}$. Thus $\beta_i = \delta = \alpha^{s-1}$ for any such i. Similarly, $(\{e_1, \ldots, e_{2t-1}\} - e_i) \cup \{f_i, z\}$ is a balanced cycle for any $i \in \{t+1, \ldots, 2t-1\}$, from which we deduce that $\beta_i = \epsilon = \alpha^s$.

As $\{f_1, \ldots, f_t, e_{t+1}, \ldots, e_{2t-1}, g\}$ and $\{e_1, \ldots, e_t, f_{t+1}, \ldots, f_{2t-1}, g\}$ are both balanced cycles in $\Delta(\mathbb{K}, t, \alpha)$, it now follows that the products $\beta_1 \cdots \beta_t = (\alpha^{s-1})^t$ and $\beta_{t+1} \cdots \beta_{2t-1} = (\alpha^s)^{t-1}$ are both equal to ζ. Thus $\alpha^{st-t} = \alpha^{st-s}$, implying $\alpha^s = \alpha^t$. Let o be the order of α in $\mathbb{L}$. Since $s \neq t$, we know that $o < \max\{s, t\}$. But if $o < s$, then $\{y_1, \ldots, y_o, x_{o+1}, \ldots, x_s, x\}$ is a balanced cycle in (G, θ), although it is not a circuit in M. Therefore $o < t$. Now the product of edge labels on the cycle $\{f_1, \ldots, f_o, e_{o+1}, \ldots, e_{2t-1}, x\}$ is $(\alpha^{s-1})^o = 1$, so this is a balanced cycle in (G, θ), although not a circuit in M'. This contradiction proves the lemma. $\square$

6. Proof of Lemma 1.4

This section is dedicated to proving Lemma 1.4, which we restate with an explicit bound. Let k be a positive integer. Define $g_2(k, 0)$ to be $2^{k^2} 3^k 7^{2k}$. Recursively define $g_2(k, n+1)$ to be $2^{g_2(k,n)}$, and let $f_2(k)$ be $g_2(k, k)$. Recall that ℓ is the set $\{a, b, x, y, z\}$, and $\mathcal{M}_\ell$ is the class of matroids having a $U_{2,5}$-restriction on ℓ. A pair of matroids is (k,ℓ)-equivalent if they have no (k,ℓ)-certificate, as defined in the introduction.

Lemma 6.1. *Let k be a positive integer. There are at most $f_2(k)$ equivalence classes of $\mathcal{M}_\ell$ under the relation of (k,ℓ)-equivalence.*

Proof. The main ideas required here are essentially identical to those in Section 3, so we omit many details. A *registry* is a $(k+2)\times k$ matrix with columns indexed by the variables $X_1,\dots,X_k$ and rows indexed by Ind, Sing, and $X_1,\dots,X_k$. As before, an entry in row X_i is either 'T' or 'F', and an entry in row Sing is either '0', '1', or '>'. Let $\mathcal{A}$ be the set

$$\{\mathrm{D},\mathrm{S}\}\cup\{\alpha\colon \alpha\subseteq\ell,\ |\alpha|\le 2\}\cup\{(\alpha,\beta)\colon \alpha,\beta\subseteq\ell,\ |\alpha|,|\beta|\le 1,\ \alpha\cap\beta=\emptyset\}.$$

A registry entry in row Ind must be a member of $\mathcal{A}$. A simple calculation shows that $|\mathcal{A}|=49$. Therefore there are at most $2^{k^2}3^k49^k=g_2(k,0)$ possible registries. A *depth-*0 *tree* is a registry, and a *depth-*$(n+1)$ *tree* is a non-empty set of depth-n trees. Hence there are no more than $f_2(k)$ depth-k trees.

A *stacked matroid* is a tuple, $\mathcal{M}=(M,Y_1,\dots,Y_m)$, where M is in $\mathcal{M}_\ell$, and each Y_i is a subset of $E(M)$. If $||\mathcal{M}||=m\le k$, then we associate a depth-$(k-||\mathcal{M}||)$ tree, $\mathcal{T}(\mathcal{M})$ to $\mathcal{M}$. We give the definition of $\mathcal{T}(\mathcal{M})$ only in the case that $\mathcal{T}(\mathcal{M})$ is a registry, because otherwise the definition is identical to that in Lemma 3.1. Assume that $\mathcal{M}=(M,Y_1,\dots,Y_k)$. The entry in row X_i and column X_j of the registry $\mathcal{T}(\mathcal{M})$ is 'T' if and only if $Y_i\subseteq Y_j$. The entry in row Sing and column X_i is '0', '1', or '>', according to whether $|Y_i|$ is less than, equal to, or greater than one.

The rules defining the entries in row Ind are more complicated. Let ω stand for the entry in row Ind and column X_j. If Y_j is dependent in M, then we set ω to be 'D'. Now we assume that Y_j is independent. Let π be the integer $r_M(Y_j-\ell)-r_M(Y_j\cup\ell)+2$. This is known as the *local connectivity* of $Y_j-\ell$ and ℓ. The submodularity of the rank function shows that $\pi\ge 0$, and since $r_M(Y_j-\ell)\le r_M(Y_j\cup\ell)$, it follows that $\pi\le 2$. If $\pi=2$, then $Y_j-\ell$ spans ℓ, and we set ω to be 'S'. In the next case, we assume that $\pi=0$. Certainly $|Y_j\cap\ell|\le 2$, as we have assumed that Y_j is independent in M. So $Y_j\cap\ell$ is in $\mathcal{A}$, and we set ω to be $Y_j\cap\ell$. Finally, we consider the case that $\pi=1$. Thus $|Y_j\cap\ell|\le 1$, for otherwise

$$r_M(Y_j)=r_M(Y_j\cup\ell)=r_M(Y_j-\ell)-\pi+2=|Y_j-\ell|+1<|Y_j|,$$

which contradicts our assumption that Y_j is independent in M. We let β be the set $Y_j\cap\ell$. Let α be $\mathrm{cl}_M(Y_j-\ell)\cap\ell$. Note that $\alpha\cap\beta=\emptyset$, as otherwise Y_j contains a circuit of M. Moreover,

$$r_M(\alpha)\le r_M(\mathrm{cl}_M(Y_j-\ell))+r_M(\ell)-r_M(\mathrm{cl}_M(Y_j-\ell)\cup\ell)=\pi=1,$$

so $|\alpha|\le 1$. Therefore (α,β) is in $\mathcal{A}$, and we set ω to be (α,β).

Let ψ be an MS_0 formula such that either ψ is quantifier-free or $\mathrm{Var}(\psi)=\{X_1,\dots,X_k\}$. Let $b(\psi)$ be the number of bound variables in ψ, and let $\mathcal{T}$ and $\mathcal{T}'$ be depth-$b(\psi)$ trees. We will define what it means for $\mathcal{T}$ and $\mathcal{T}'$ to be ψ-compatible. We give the definition only in the case that $b(\psi)=0$ and ψ is the atomic formula $\mathrm{Ind}(X_j)$: otherwise the definition is identical to that in Lemma 3.1. Let ω and ω' be the entries of $\mathcal{T}$ and $\mathcal{T}'$ in row Ind and column X_j. It is easiest to define the rules that determine the ψ-compatibility of $\mathcal{T}$ and $\mathcal{T}'$ via a flowchart, which is exactly what we do in Figure 3. When following this flowchart, we start in the shaded cell.

A terminal node that is hollow signifies that $\mathcal{T}$ and $\mathcal{T}'$ are ψ-compatible. A filled terminal node signifies that they are not. Note that if ω is not 'D' or 'S', then it is either a subset of ℓ or a pair (α, β), where α and β are subsets of ℓ. The same comment applies to ω'.

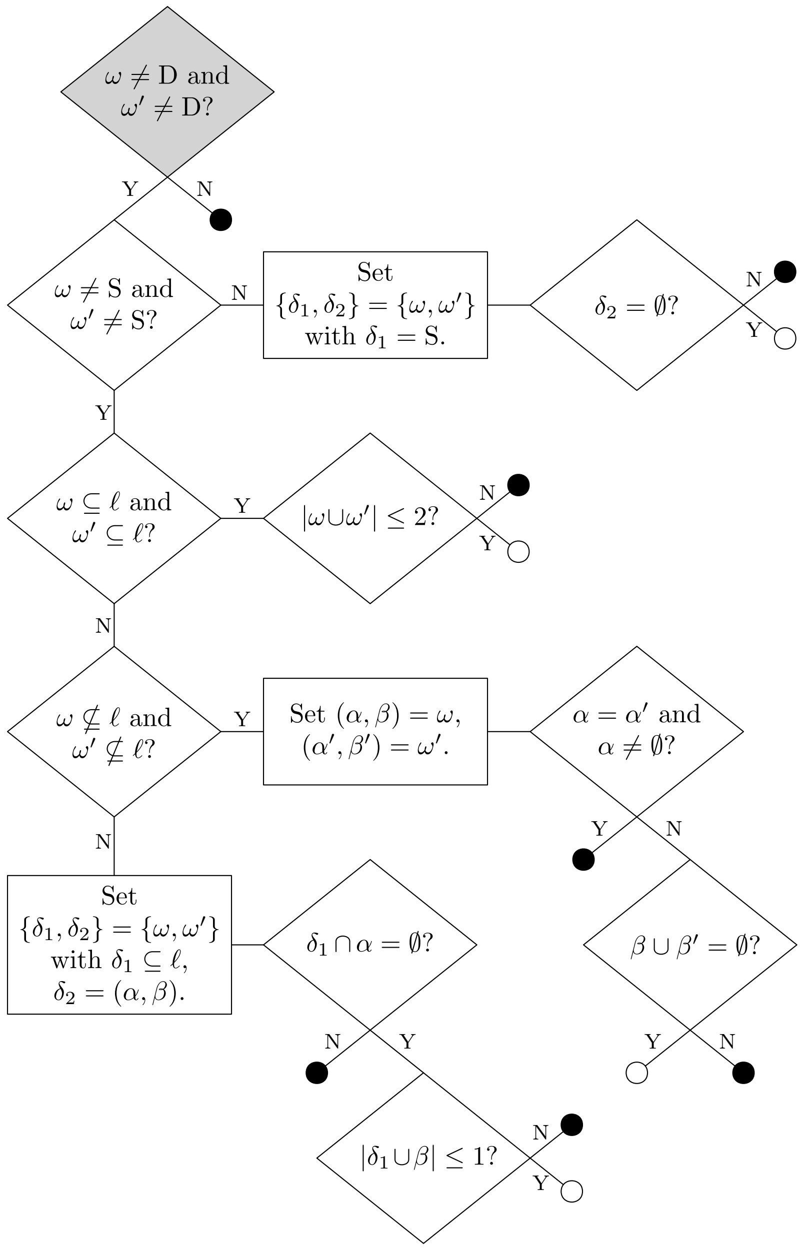

FIGURE 3. Deciding whether $\mathcal{T}$ and $\mathcal{T}'$ are ψ-compatible.

Claim 6.1.1. Let ψ be an MS_0 formula such that either ψ is quantifier-free or $\mathrm{Var}(\psi) = \{X_1, \dots, X_k\}$. If $\mathrm{Var}(\psi) = \{X_1, \dots, X_k\}$, then let m be $|\mathrm{Fr}(\psi)|$ and assume that $\mathrm{Fr}(\psi) = \{X_1, \dots, X_m\}$. Otherwise, let m be k. Let M and M' be matroids in $\mathcal{M}_\ell$ satisfying $E(M) \cap E(M') = \ell$, and let $\mathcal{M} = (M, Y_1, \dots, Y_m)$ and $\mathcal{M}' = (M', Y_1', \dots, Y_m')$ be stacked matroids. Define τ to be the function that takes X_i to $Y_i \cup Y_i'$, for each $X_i \in \mathrm{Fr}(\psi)$. The interpretation $(\mathrm{Amal}(M, M'), \tau)$ satisfies ψ if and only if the trees $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible.

Proof. The proof of this claim differs from that of Claim 3.1.1 only in the base case when ψ is the atomic formula $\mathrm{Ind}(X_j)$. Therefore we need only consider this case. Let ψ be the formula $\mathrm{Ind}(X_j)$. Let ω be the entry in row Ind and column X_j of the registry $\mathcal{T}(\mathcal{M})$, and let ω' be the corresponding entry of $\mathcal{T}(\mathcal{M}')$. We will trace all possible outcomes in the flowchart shown in Figure 3. We will prove that if $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible, then $Y_j \cup Y_j'$ is independent in $\mathrm{Amal}(M, M')$, whereas if they are not ψ-compatible, then $Y_j \cup Y_j'$ is dependent. This will establish the claim. Let X be the set $Y_j \cup Y_j'$.

If either ω or ω' is 'D', then either Y_j is dependent in M or Y_j' is dependent in M'. In this case $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are not ψ-compatible, and X is certainly dependent in $\mathrm{Amal}(M, M')$. Therefore we will assume that $\omega \neq \mathrm{D}$ and $\omega' \neq \mathrm{D}$, so Y_j is independent in M and Y_j' is independent in M'.

In the next case, we assume that either ω or ω' is 'S'. By symmetry, we can assume that $\omega = \mathrm{S}$. Then $Y_j - \ell$ spans ℓ in M. Since Y_j is independent in M, we observe that $Y_j - \ell = Y_j$. Assume that $\omega' \neq \emptyset$, so that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are not ψ-compatible. If ω' is a non-empty subset of ℓ, then $\omega' = Y_j' \cap \ell$, and it follows that an element of Y_j' is in the closure of $Y_j - \ell$ in M, so that X is dependent. If ω' is not a subset of ℓ, then $r_{M'}(Y_j' - \ell) - r_{M'}(Y_j' \cup \ell) + 2 > 0$, meaning that $r_{M'}((X - E(M)) \cup \ell) < r_{M'}(X - E(M)) + 2$. Thus Proposition 4.1 implies that X is dependent in $\mathrm{Amal}(M, M')$. On the other hand, if $\omega' = \emptyset$, then $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible. Furthermore, $r_{M'}(Y_j' - \ell) - r_{M'}(Y_j' \cup \ell) + 2 = 0$ and $Y_j' \cap \ell = \emptyset$, meaning that $Y_j' - \ell = Y_j'$. Now we know that $X \cap \ell = \emptyset$, so that $X \cap E(M)$ is independent in M and $X \cap E(M')$ is independent in M'. The fact that $r_{M'}(Y_j') + 2 = r_{M'}(Y_j' \cup \ell)$ implies that $\mathrm{cl}_{M'}(Y_j') \cap \ell = \emptyset$. None of the statements in Proposition 4.1 apply, so X is independent in $\mathrm{Amal}(M, M')$.

We now follow the branch of the flowchart in which $\omega \neq \mathrm{S}$ and $\omega' \neq \mathrm{S}$. This means that neither $Y_j - \ell$ nor $Y_j' - \ell$ spans ℓ. Assume that ω and ω' are both subsets of ℓ. This implies that $r_M(Y_j - \ell) - r_M(Y_j \cup \ell) + 2$ and $r_{M'}(Y_j' - \ell) - r_{M'}(Y_j' \cup \ell) + 2$ are both zero. From this we deduce that $\mathrm{cl}_M(Y_j - \ell) \cap \ell$ and $\mathrm{cl}_{M'}(Y_j' - \ell) \cap \ell$ are empty. Assume that $|\omega \cup \omega'| > 2$. Then $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are not ψ-compatible. As ℓ is a rank-2 set, obviously it follows that $X \cap E(M)$ and $X \cap E(M')$ are dependent. Therefore we assume that $|\omega \cup \omega'| \leq 2$, so that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible. As $r_M(Y_j - \ell) = r_M(Y_j \cup \ell) - 2$, and $X \cap \ell = \omega \cup \omega'$ contains at most two elements, we see that $X \cap E(M)$ is independent in M. By exactly the same argument, $X \cap E(M')$ is independent in M'. The information we have assembled in this paragraph is enough to determine that none of the statements in Proposition 4.1 apply, so X is independent in $\mathrm{Amal}(M, M')$.

Next we consider the branch where neither ω nor ω' is a subset of ℓ. This means that both $r_M(Y_j - \ell) - r_M(Y_j \cup \ell) + 2$ and $r_{M'}(Y_j' - \ell) - r_{M'}(Y_j' \cup \ell) + 2$ are equal to one. Let ω be (α, β), where α and β are disjoint subsets of ℓ of size at most one, and similarly assume that $\omega' = (\alpha', \beta')$. Assume that $\alpha = \alpha'$ and that $\alpha \neq \emptyset$, so that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are not ψ-compatible. The single element in α belongs to both $\mathrm{cl}_M(Y_j - \ell)$ and $\mathrm{cl}_{M'}(Y'_j - \ell)$. Statement (iii) of Proposition 4.1 now implies that X is dependent. Thus we assume that either $\alpha \neq \alpha'$ or $\alpha = \alpha' = \emptyset$. Assume that $\beta \cup \beta' \neq \emptyset$, so that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are not ψ-compatible. By symmetry, we will assume that $\beta \neq \emptyset$ and e is the single element in β. Then e is in $Y_j \cap \ell$, but not in $\mathrm{cl}_M(Y_j - \ell)$. Since $r_M(Y_j \cup \ell) = r_M(Y_j - \ell) + 1$, we now see that Y_j spans ℓ in M. As $r_{M'}(Y'_j \cup \ell) = r_{M'}(Y_j - \ell) + 1$, Proposition 4.1 tells us that X is dependent. On the other hand, if $\beta \cup \beta = \emptyset$, then $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible and $X \cap \ell$ is empty, which means that $X \cap E(M)$ is independent in M and $X \cap E(M')$ is independent in M'. Earlier we followed the branch in which neither $\mathrm{cl}_M(Y_j - \ell)$ nor $\mathrm{cl}_M(Y'_j - \ell)$ contains ℓ. It follows that neither $\mathrm{cl}_M(X \cap E(M))$ nor $\mathrm{cl}_{M'}(X \cap E(M'))$ contains ℓ. There is no element of ℓ in both $\mathrm{cl}_M(Y_j - \ell)$ nor $\mathrm{cl}_M(Y'_j - \ell)$, since in that case the element would be in α and α'. Therefore Proposition 4.1 implies that X is independent.

Finally we arrive at the branch of the flowchart where exactly one of ω and ω' is a subset of ℓ. By symmetry, we will assume that $\omega' \subseteq \ell$ and $\omega = (\alpha, \beta)$, where α and β are disjoint subsets of ℓ of size at most one. If there is an element of ω' in α, then this element is in $(Y'_j \cap \ell) \cap \mathrm{cl}_M(Y_j - \ell)$, which implies that $X \cap E(M)$ is dependent in M. As $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are not ψ-compatible in this branch, this is the desired outcome. Therefore we assume that $\omega' \cap \alpha = \emptyset$. Assume that $\omega' \cup \beta$ contains distinct elements e and f. This means that $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are not ψ-compatible. We have just assumed that $\omega' \cap \alpha = \emptyset$, from which it follows that e is not in $\mathrm{cl}_M(Y_j - \ell)$. As $r_M(Y_j - \ell) = r_M(Y_j \cup \ell) - 1$, we deduce that $r_M((Y_j - \ell) \cup e) = r_M(Y_j \cup \ell)$. Therefore $(Y_j - \ell) \cup e$ spans f in M, so $X \cap E(M)$ is dependent. Now we assume that $\omega' \cup \beta$ contains at most one element. Therefore $\mathcal{T}(\mathcal{M})$ and $\mathcal{T}(\mathcal{M}')$ are ψ-compatible. Since $\omega' \cap \alpha = \emptyset$, it follows easily that $(Y_j - \ell) \cup (\omega' \cup \beta) = X \cap E(M)$ is independent in M. Similarly, $X \cap E(M') = (Y'_j - \ell) \cup (\omega' \cup \beta)$ is independent in M'. Because $r_{M'}(Y'_j - \ell) - r_{M'}(Y'_j \cup \ell) + 2 = 0$, there is no element in $\mathrm{cl}_{M'}(Y'_j - \ell) \cap \ell$. Proposition 4.1 implies that the only way X can be dependent in $\mathrm{Amal}(M, M')$ is if ℓ is contained in $\mathrm{cl}_{M'}(X \cap E(M'))$. But this is impossible, as $r_{M'}(Y_j - \ell) = r_{M'}(Y'_j \cup \ell) - 2$, and there is at most one element in $X \cap \ell$. Therefore X is independent in $\mathrm{Amal}(M, M')$, exactly as desired. □

We complete the proof of Lemma 6.1 by observing that Claim 6.1.1 implies that the number of (k, ℓ)-equivalence classes is bounded above by the number of depth-k trees. □

We can now prove Theorem 1.2 and Corollaries 1.5 and 1.6.

Proof of Theorem 1.2. Let $\mathbb{K}$ be an infinite field. Assume that $\psi_{\mathbb{K}}$ is a sentence in MS_0 characterising $\mathbb{K}$-representable matroids. Observe that $\mathbb{K}$ contains non-zero elements with arbitrarily large order: to see this, assume that the order of every element in $\mathbb{K} - \{0\}$ is bounded above by the integer K. Then every element in $\mathbb{K} - \{0\}$ is a root of the polynomial $(x^K - 1)(x^{K-1} - 1) \cdots (x - 1)$. Since there are only finitely many such roots, $\mathbb{K}$ is finite. This contradiction proves our claim.

Let k be $|\mathrm{Var}(\psi_{\mathbb{K}})|$. We apply Lemma 6.1. Choose the element $\alpha \in \mathbb{K} - \{0\}$ with high enough order so that there are at least $f_2(k) + 1$ integers s such that $s \geq 3$ and $2s(s-1)$ is less than the order of α. Then there are two distinct integers, s and t, satisfying these constraints, such that $M_1 = M(\Gamma(\mathbb{K}, s, \alpha))$ and $M_2 = M(\Gamma(\mathbb{K}, t, \alpha))$

are (k,ℓ)-equivalent. We let M' be $M(\Delta(\mathbb{K},s,\alpha))$. Then $\psi_{\mathbb{K}}$ is satisfied by both of $\mathrm{Amal}(M_1,M')$ and $\mathrm{Amal}(M_2,M')$ or by neither. However, the first of these amalgams is $\mathbb{K}$-representable by Lemma 5.2, and the second is not representable over any field at all by Lemma 5.3. This contradiction completes the proof of the theorem. $\square$

Proof of Corollary 1.5. Let $\{\psi_q\}_{q\in\mathcal{Q}}$ be a set of sentences characterising $\mathrm{GF}(q)$-representability, and assume that N is an integer such that $|\mathrm{Var}(\psi_q)| \leq N$ for all $q \in \mathcal{Q}$. Recall that if $q \in \mathcal{Q}$, then the multiplicative group of $\mathrm{GF}(q)$ has an element of order $q-1$. We apply Lemma 6.1. Choose $q \in \mathcal{Q}$ large enough so that there are at least $f_2(N)+1$ integers s satisfying $s \geq 3$ and $2s(s-1) < q-1$. Let α be a generator of the multiplicative group of $\mathrm{GF}(q)$. Assume that ψ_q contains $k \leq N$ variables. As $f_2(N)+1 \geq f_2(k)+1$, there are distinct integers, s and t, such that $s,t \geq 3$ and $2s(s-1), 2t(t-1) < q-1$ and $M = M(\Gamma(\mathbb{K},s,\alpha))$ and $M' = M(\Gamma(\mathbb{K},t,\alpha))$ are (k,ℓ)-equivalent. Now we obtain a contradiction from Lemmas 5.2 and 5.3 exactly as before. $\square$

Proof of Corollary 1.6. If $\mathbb{K}$ is an infinite field with characteristic c, then $\mathbb{K}$ contains elements of arbitrarily high order, so all matroids of the form $M_1 = M(\Gamma(\mathbb{K},s,\alpha))$ and $M_2 = M(\Gamma(\mathbb{K},t,\alpha))$ are $\mathbb{K}$-representable. Therefore the proof proceeds exactly as in Theorem 1.2. $\square$

Acknowledgments

We thank Noam Greenberg for helpful advice and the referees for their constructive feedback. The research in this article was supported by the Rutherford Discovery Fellowship, the Marsden Fund of New Zealand, and NSERC Canada.

References

[1] Martin Aigner, *Combinatorial theory*, reprint of the 1979 original, Classics in Mathematics, Springer-Verlag, Berlin, 1997. MR1434477

[2] Bruno Courcelle, *The monadic second-order logic of graphs. I. Recognizable sets of finite graphs*, Inform. and Comput. **85** (1990), no. 1, 12–75. MR1042649

[3] Heinz-Dieter Ebbinghaus and Jörg Flum, *Finite model theory*, Second revised and enlarged edition, Springer Monographs in Mathematics, Springer-Verlag, Berlin, 2006. MR2177676

[4] Jim Geelen, *Some open problems on excluding a uniform matroid*, Adv. in Appl. Math. **41** (2008), no. 4, 628–637. MR2459453

[5] Jim Geelen, Bert Gerards, and Geoff Whittle, *Solving Rota's conjecture*, Notices Amer. Math. Soc. **61** (2014), no. 7, 736–743. MR3221124

[6] Petr Hliněný, *On matroid properties definable in the MSO logic*, Mathematical foundations of computer science 2003, Lecture Notes in Comput. Sci., vol. 2747, Springer, Berlin, 2003, pp. 470–479. MR2081597

[7] Thomas Lengauer and Egon Wanke, *Efficient analysis of graph properties on context-free graph languages (extended abstract)*, Automata, languages and programming (Tampere, 1988), Lecture Notes in Comput. Sci., vol. 317, Springer, Berlin, 1988, pp. 379–393. MR1023649

[8] Leonid Libkin, *Elements of finite model theory*, Texts in Theoretical Computer Science. An EATCS Series, Springer-Verlag, Berlin, 2004. MR2102513

[9] Dillon Mayhew, Geoff Whittle, and Mike Newman, *Is the missing axiom of matroid theory lost forever?*, Q. J. Math. **65** (2014), no. 4, 1397–1415. MR3285777

[10] A. Nerode, *Linear automaton transformations*, Proc. Amer. Math. Soc. **9** (1958), 541–544. MR0135681

[11] James Oxley, *Matroid theory*, 2nd ed., Oxford Graduate Texts in Mathematics, vol. 21, Oxford University Press, Oxford, 2011. MR2849819

[12] P. Vámos, *The missing axiom of matroid theory is lost forever*, J. London Math. Soc. (2) **18** (1978), no. 3, 403–408. MR518224

[13] Hassler Whitney, *On the abstract properties of linear dependence*, Amer. J. Math. **57** (1935), no. 3, 509–533. MR1507091

[14] Thomas Zaslavsky, *Biased graphs. IV. Geometrical realizations*, J. Combin. Theory Ser. B **89** (2003), no. 2, 231–297. MR2017726

School of Mathematics and Statistics, Victoria University of Wellington, P.O. Box 600, Wellington 6140, New Zealand

Email address: `dillon.mayhew@vuw.ac.nz`

Department of Mathematics and Statistics, University of Ottawa, Ottawa, Ontario K1N 6N5, Canada

Email address: `mnewman@uottawa.ca`

School of Mathematics and Statistics, Victoria University of Wellington, P.O. Box 600, Wellington 6140, New Zealand

Email address: `geoff.whittle@vuw.ac.nz`