

Disasters in Abstracting Combinatorial Properties of Linear Dependence

by

Rutger Theodoor Ronald Jansen van Doorn Campbell

A thesis
presented to the University of Waterloo
in fulfillment of the
thesis requirement for the degree of
Doctor of Philosophy
in
Combinatorics and Optimization

Waterloo, Ontario, Canada, 2020

© Rutger Theodoor Ronald Jansen van Doorn Campbell 2020

Examining Committee Membership

The following served on the Examining Committee for this thesis. The decision of the Examining Committee is by majority vote.

External Examiner: Mike Newman
Associate Professor, Dept. of Mathematics and Statistics,
University of Ottawa

Supervisor: Jim Geelen
Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Internal Member: Peter Nelson
Assistant Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Internal Member: David Wagner
Professor, Dept. of Combinatorics and Optimization,
University of Waterloo

Internal-External Member: Ross Willard
Professor, Dept. of Pure Mathematics,
University of Waterloo

Author's Declaration

I hereby declare that I am the sole author of this thesis. This is a true copy of the thesis, including any required final revisions, as accepted by my examiners.

I understand that my thesis may be made electronically available to the public.

Abstract

A notion of geometric structure can be given to a set of points without using a coordinate system by instead describing geometric relations between finite combinations of elements. The fundamental problem is to then characterize when the points of such a “geometry” have a consistent coordinatization. Matroids are a first step in such a characterization as they require that geometric relations satisfy inherent abstract properties.

Concretely, let E be a finite set and \mathcal{I} be a collection of subsets of E . The problem is to characterize pairs (E, \mathcal{I}) for which there exists a “representation” of E as vectors in a vector space over a field \mathbb{F} where \mathcal{I} corresponds to the linear independent subsets of E . Necessary conditions for such a representation to exist include: the empty set is independent, subsets of independent sets are also independent, and for each subset X , the maximal independent subsets of X have the same size. When these properties hold, we say that (E, \mathcal{I}) describes a *matroid*. As a result of these properties, matroids provide many useful concepts and are an appropriate context in which to consider characterizations.

Mayhew, Newman, and Whittle showed that there exist pathological obstructions to natural axiomatic and forbidden-substructure characterizations of real-representable matroids. Furthermore, an extension of a result of Seymour illustrates that there is high computational complexity in verifying that a representation exists. This thesis shows that such pathologies still persist even if it is known that there exists a coordinatization with complex numbers and a sense of orientation, both of which are necessary to have a coordinatization over the reals.

Acknowledgements

I have had the pleasure of having Jim Geelen as my supervisor. He has been a great mentor and has provided many excellent problems and answers. I thank him for his insights both academic and professional and for his patience. He has helped me distinguish the forest for the trees (as long as they are not the mathematical objects).

I thank my examining committee members, Peter Nelson, Mike Newman, Ross Willard, and David Wagner for their careful reading, helpful comments, and interesting questions. I also thank Peter for his feedback and perspective over the years, and for his good humour in our research together.

I thank James Oxley and Geoff Whittle, for their insights into math and language over the years and in their many proofreadings of this thesis.

I thank Anneke Jansen van Doorn, Michael Campbell, Samuel Baltz, and Chris Hawthorne for proofreading various sections and helpful discussions. I thank them and other family and friends for all their support and keeping me sane.

This research was partially supported by a scholarship from the Natural Sciences and Engineering Research Council of Canada.

Dedication

To JavaDoCa.

And to Michael and Anneke, who first taught me math and language in different combinations — including my first exposure to presenting geometric constructions.

Table of Contents

1	The fundamental question of matroid theory	1
1.1	Formalism	2
1.2	Axiomatic characterization	3
1.3	Forbidden-substructure characterization	4
1.4	Algorithmic characterization	4
1.5	Orientable matroids	5
2	Encoding algebra in matroids	7
2.0.1	Representations	8
2.0.2	Parametrization	8
2.0.3	Encoding values as points on lines	9
2.0.4	Algebraic relations through restrictions	10
2.1	Local algebraic structure; von Staudt constructions	10
2.1.1	Non-degenerate addition	11
2.1.2	Non-degenerate multiplication	12
2.1.3	Squaring/Inversion	13
2.1.4	Algebraic extensions	14
2.2	Global algebraic structure	15
2.2.1	Sums; Spikes	15
2.2.2	Products; Swirls	17
2.3	Sharing structure; Amalgams	19
2.3.1	Cannibalizing structure; Pinned extensions	22

3	Axiomatization	25
3.0.1	The independent-set language MS_0	25
3.1	Matroid axiomatization	26
3.1.1	Important sets in a matroid	27
3.1.2	Minors	28
3.1.3	Ingleton’s inequality	29
3.2	Ineffable properties	30
3.2.1	Depth- k truth tables	30
3.2.2	k -equivalence	31
3.3	The non-axiomatizability of representability	32
3.4	The non-axiomatizability of real-representability	33
3.4.1	Alterations to the previous technique	33
3.4.2	New constructions	34
3.4.3	Analogous proof	35
4	Excluded minors	36
4.0.1	Excluded minors for real-representability	36
4.0.2	Complex-representability preserving operations	37
4.1	Natural classes	38
4.1.1	Gammoids	39
4.2	Constructing a “major” excluded minor	41
4.2.1	Picking a base excluded minor N in \mathcal{N}	41
4.2.2	Preprocessing M ; Bipartition by bases	42
4.2.3	Using M to cannibalize N	42
5	Computational Complexity	48
5.1	Decidability; solvability of systems of integer polynomial equations	49
5.2	Certification	49

5.2.1	Certifying minors	50
5.2.2	Not polynomially certifiable	50
5.3	Real-representability is not polynomially certifiable	51
5.4	Certifying non-representability for finite fields	52
5.5	Non-real-representability is not polynomially certifiable	53
5.5.1	A restricted representation problem	53
5.5.2	Enforcing this restricted problem	55
5.6	Assuming complex-representability	59
5.6.1	Real-representability is still not polynomially certifiable	59
5.6.2	Non-real-representability is still not polynomially certifiable	61
6	Orientability	62
6.0.1	Orienting circuits	63
6.0.2	Orienting the complements of hyperplanes	65
6.1	Orientability-preserving operations	66
6.1.1	Circuit-hyperplane relaxation	66
6.1.2	Amalgams of orientable matroids	67
6.2	Representable orientable matroids that are not real-representable	68
6.3	Which constructions still work?	71
	REFERENCES	73

The fundamental question of completely characterizing systems which represent matrices is left unsolved.

— Hassler Whitney, 1935

Chapter 1

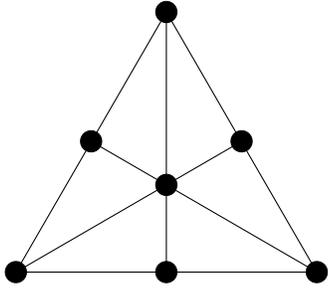
The fundamental question of matroid theory

How easily can we describe points from a Euclidean space without using coordinates? An appropriate context for this question is matroid theory, as it provides a convenient geometric framework with many useful concepts [34]. Matroids were first introduced in 1935 by Nakasawa and by Whitney to abstract the geometric structure of a finite collection of vectors in a vector space [31,45]. Once we forget the underlying vector space, the resulting geometric structure is that of a matroid (which will be properly defined later).

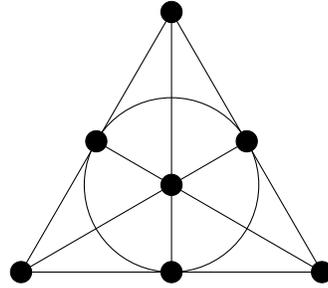
Any axiomatization of matroids arises from fundamental conditions on collections of vectors. However, Whitney gave an example of a matroid that does not come from any collection of vectors in a real vector space (see Figure 1.1). Whitney asked for a characterization of “Euclidean” matroids, or equivalently, those that can be represented as a collection of vectors from a real vector space. So far any attempt to usefully characterize real-representability has failed spectacularly [27,28]. This thesis will prove several results that reinforce how intractable this question is.

Real-representability can potentially be characterized algorithmically, by adding extra axioms, or by describing structural obstructions. While each of these approaches will be elaborated upon later, any natural implementation is found wanting. A useful characterization of real-representability would need to use conditions that are comparably complicated to real-representability in each of these settings.

For any field \mathbb{F} , having an \mathbb{F} -representation is a strong property quantitatively. Matroids in general are incredibly wild in comparison to the class of matroids that can be represented over at least one field. The number of n -element matroids is $2^{\frac{1}{\text{poly}(n)}}2^n$ whereas the number



(i) A real-representable matroid (as it has a straight line drawing).



(ii) The non-real-representable matroid given by Whitney (see Theorem 5.1.2 for an algebraic technique to show this).

Figure 1.1: Although very similar to each other, only (i) is real-representable, while (ii) is minimally non-real-representable.

of n -element representable matroids is only $2^{\text{poly}(n)}$, where, in each case, $\text{poly}(n)$ is some function that is bounded above and below by polynomials; see Knuth [20] and Nelson [32], respectively.

A representation with complex vectors is an immediate necessary condition for real-representability. Complex-representability is also found to be difficult to describe in the aforementioned settings [27,28]. This would lead one to have strong hopes for characterizing real-representability for the complex-representable matroids. Remarkably, however, this simplification makes no discernible difference from the perspective of the methodologies alluded to above. Algebraic considerations are found to still play a large role even when the geometric structure is well-behaved.

1.1 Formalism

We will now go into further detail on the characterization techniques considered. First, however, we see concrete ways we could describe the structure of a matroid.

Imagine that you are given a pair (E, \mathcal{I}) consisting of a *ground set* E together with a collection \mathcal{I} of *independent* subsets of E . Consider whether \mathcal{I} relates the linearly independent subsets of some collection of vectors. Is there a multiset, $\{v_e\}_{e \in E}$, of vectors in a vector space over a field \mathbb{F} such that \mathcal{I} corresponds to the linearly independent subsets of $\{v_e\}_{e \in E}$? Some of the elementary conditions that such an \mathcal{I} must satisfy give rise to the definition of a matroid. Specifically: the empty set is independent, subsets of independent

sets are also independent, and each for each subset, X , the maximal independent subsets of X have the same size (called the *rank of X* , denoted $r(X)$). When these properties hold, we call $M = (E, \mathcal{I})$ a *matroid*.

We can get alternate axiomatizations of matroids with equivalent conditions if we instead use different concepts to capture the geometric structure. For instance, the maximal independent sets (the *bases*), or the maximal sets not containing a basis (the *hyperplanes*), both of which are analogous to their eponyms for vector spaces.

We say that $M = (E, \mathcal{I})$ is an \mathbb{F} -*representable* or *representable matroid* when there actually exists a multiset $\{v_e\}_{e \in E}$ of \mathbb{F} -vectors that has the same set of independent subsets, or equivalently, has the same structure as M in terms of any of the notions just mentioned. We now consider possible characterizations of real-representable matroids.

1.2 Axiomatic characterization

Consider whether we can define real-representability by adding more constraints to \mathcal{I} in the same language as we used to define matroids in the first place. We already know that we can define representable matroids if we have a strong enough logical language: we have already done so informally above. In Section 3.0.1, we will give a natural independent-set language, MS_0 , first developed by Mayhew, Newman, and Whittle [27] and equivalent to one used by Hliněný [15]. With this independent-set language in mind, Mayhew, Newman, and Whittle [27] showed that real-representability is ineffable.

[Theorem ?? (Mayhew, Newman, Whittle [27])]. *Then is no sentence $\phi_{\mathbb{R}}$ in MS_0 such that a matroid is real-representable precisely when it satisfies $\phi_{\mathbb{R}}$.*

They further showed that neither representability nor complex-representability is finitely axiomatizable in MS_0 . Square brackets were used around the above theorem and are used in general to indicate where the corresponding statement occurs most naturally.

One might hope that the difficulty lies in defining representability, and once we have a characterization of representability, it is an easier matter to define representability over a specific field. However, knowing that a matroid is representable or even complex-representable does not lead to a finite condition (in MS_0) for real-representability.

[Theorem 3.4.2 (Campbell)]. *There is no sentence ϕ in MS_0 such that a complex-representable matroid is real-representable precisely when it satisfies ϕ .*

1.3 Forbidden-substructure characterization

We now consider characterizing obstructions to real-representability. Given a matroid $M = (E, \mathcal{I})$ and an element $e \in E$, there are two natural methods to derive a new matroid $M' = (E - \{e\}, \mathcal{I}')$. We can take \mathcal{I}' to be all the sets in \mathcal{I} that do not contain e (*deletion of e* , denoted $M \setminus e$). Alternatively, we can “project” from e when $\{e\} \in \mathcal{I}$ by taking \mathcal{I}' to be all the sets $I - e$ where $e \in I \in \mathcal{I}$ (*contraction of e* , denoted M/e). When there is no set in \mathcal{I} containing e , we define $M/e = M \setminus e$. Intuitively, these operations get rid of some of the structure of M and create a “smaller” matroid. When a matroid N can be obtained from M with a sequence of these operations, we call N a *minor* of M . If we only used deletion to obtain N on ground set $S \subseteq E$, we say that N is the *restriction* of M to S , denoted $M|S$.

The class of real-representable matroids is closed under taking minors. This means that we can characterize real-representable matroids by giving the minor-minimal matroids that are not real-representable — the *excluded minors*. However, Mayhew, Newman, and Whittle [28] showed that, remarkably, the excluded minors are at least as wild as the class of real-representable matroids themselves. More precisely, the following is a special case of Theorem 4.0.1.

Theorem 1.3.1 (Mayhew, Newman, Whittle [28]). *Each real-representable matroid is a minor of an excluded minor for real-representability.*

This is not as surprising when we recall the quantitative comparison between matroids and representable matroids: almost all matroids are non-representable [32]. However, the class of obstructions for real-representability remains intractable even when restricting to those that are complex-representable.

[**Theorem 4.0.2** (Campbell, Geelen [8])]. *Each real-representable matroid is a minor of a complex-representable excluded minor for real-representability.*

1.4 Algorithmic characterization

What can one say about the representability of a matroid and how easily? Mercifully, real-representability is decidable: we can characterize real-representability algorithmically (folklore with quantifier elimination [40]; see Section 5.1). However, we should consider the required complexity of such an algorithm. Concretely, how much of the matroid’s structure

needs to be queried to determine whether or not the matroid is real-representable? If we already know the answer, how much structure do we need to prove this to another party? This will be measured as the size of the ground set grows; we will make this precise in Section 5.2. By applying a technique of Seymour [36], we will see that we cannot always show that a matroid is real-representable with a polynomial number of queries of \mathcal{I} . To even certify that a matroid is **not** real-representable may require more than a polynomial number of queries in the worst case.

Certification does not become any more feasible in computational complexity even if we have prior knowledge that the given matroid is complex-representable.

[Theorem 5.0.1 (Campbell)]. *Real-representability is not polynomially certifiable even within the class of complex-representable matroids.*

[Theorem 5.6.2 (Campbell)]. *Non-real-representability is not polynomially certifiable even within the class of complex-representable matroids.*

1.5 Orientable matroids

Geometric structure can be further constrained by considering matroid “orientations”. These impose a relative sense of direction in the structure given by a matroid. The axioms governing matroid orientations arise naturally when considering collections of vectors over an ordered field. Specifically, they give “signs” in dependencies and “sides” to the complements of hyperplanes. Notably, every real-representable matroid is orientable. Conversely, there are non-orientable matroids that are representable over some field, and orientable matroids that are not representable over any field [2]. However, Whittle suggested that together these two necessary conditions for real-representability may be sufficient (personal communication, 2017):

[Conjecture 6.0.1 (Whittle)]. *A matroid is real-representable if and only if it is orientable and representable over some field.*

We will see that this conjecture is false. Indeed,

[Theorem 6.2.1 (Campbell)]. *For every finite field \mathbb{F} with $|\mathbb{F}^*| = |\mathbb{F}| - 1$ composite, there is an \mathbb{F} -representable, complex-representable, orientable matroid that is not real-representable.*

In other words, for each prime power $q \geq 5$, if q is not 2^n for some integer n where $2^n - 1$ is a (Marsenne) prime, then there is a $\text{GF}(q)$ -representable counterexample to Conjecture 6.0.1, above, that is also complex-representable.

Finally, we see that, even with orientability and complex-representability, it still seems impractical to characterize real-representability algorithmically, by adding axioms, or by forbidding substructures.

While all of these negative results rule out many possible characterizations, they do keep us realistic in our expectations of structure related to representability over infinite fields; we now know where to look. For instance, we may still hope that an alteration to Whittle's conjecture holds.

[Conjecture 6.3.5 (Revision of Whittle's conjecture)]. *If an orientable matroid is representable over some field, then it is complex-representable.*

Chapter 2

Encoding algebra in matroids

Matroids are defined by using some of the most natural properties of linear independence. However, there are other geometric conditions a matroid must satisfy to be representable. One of the most basic is Ingleton's inequality, which is an inclusion-exclusion principle on the dimensions of subspaces (see Section 3.1.3). This necessary condition for representability and others like it, see [19], can be violated to cause pathologies in some of the basic characterizations discussed in the previous chapter, see [7,28,36]). However, if we wish to show that distinguishing real-representability is complicated within the class of representable matroids, we do not have the liberty to use such conditions. Instead, the matroid obstructions must be algebraic in nature, relying on algebraic conditions that are not possible in the reals but that can be satisfied in other fields. To construct such obstructions, we need to be able to encode these algebraic conditions geometrically.

The idea of using geometry to encode algebra dates back to the ancient Greeks. However, Greek constructions rely on lengths, a concept we do not have in matroids. Constructions that only use points and lines were first developed in 1857 by von Staudt in what he called an "algebra of throws" [38]. MacLane published the first matroid theory application of these in 1936, where he gave an example of a real-representable matroid that is not representable over the rationals [24]. Since then, there have been many results where von Staudt constructions are used to give examples of matroids that are only representable when certain algebraic conditions are met, see [39,5]. Of particular note, Mnëv's Universality Theorem further considers all real-representations that are possible for a given "oriented" matroid, see [29].

2.0.1 Representations

For a representable matroid M , we may consider a representation $\{v_e\}_{e \in E(M)}$ of M as a function that takes $e \in E(M)$ to v_e or as a matrix where v_e is the column indexed by $e \in E(M)$. Here, we consider the index sets of rows and columns as unordered, but pick some ordering with which to write the matrix. We will only compute determinants to identify singular submatrices, so the sign of the determinant and hence the order of the index sets do not matter.

We say that two matrices are *row equivalent* when one can be obtained from the other by elementary row operations. We say that two matrices are *projectively equivalent* when one can be obtained from the other by elementary row operations **and** column scaling. When comparing matrices with a different number of rows we will adjoin zero rows as necessary as this does not change the row space. While row equivalent matrices are projectively equivalent, and projectively equivalent matrices represent isomorphic matroids, the converses do not hold in general. However, certain matrix representations will be more convenient, so we will often consider matrices up to some equivalence.

2.0.2 Parametrization

As representable-matroid constructions are often not dependent on the representation or the field, we will parametrize them. Representations will be over a field that is generated by a set X of indeterminates and a set Q of irreducible integer polynomials that we equate with zero (and cannot divide by). This is the field of fractions of the quotient polynomial ring $\mathbb{Z}[X]/\langle Q \rangle$. Here, $\mathbb{Z}[X]$ denotes integer polynomials in X and $\langle Q \rangle$ denotes the ideal generated by $Q \subseteq \mathbb{Z}[X]$. We will typically assume that Q is empty unless otherwise required by the algebraic relations that exist.

We say that $\alpha_1, \dots, \alpha_k$ in a field \mathbb{F} are *n -algebraically independent modulo Q* when, for any integer polynomial f of degree at most n , we have $f(\alpha_1, \dots, \alpha_k) = 0$ if and only if $f \in \langle Q \rangle$. We observe the following:

Lemma 2.0.1. *Let M be a matroid with matrix representation $A(X)$ over $\mathbb{Z}[X]/\langle Q \rangle$ where the determinants of submatrices of $A(X)$ have degree at most n . Let f be a ring homomorphism from $\mathbb{Z}[X]/\langle Q \rangle$ to a field \mathbb{F} which maps elements of X to distinct values $\alpha_1, \dots, \alpha_{|X|}$ that are n -algebraically independent modulo Q . The matrix, $A(\alpha_1, \dots, \alpha_{|X|})$, obtained by applying f to every entry of $A(X)$, is an \mathbb{F} -representation of M .*

Proof. Let S be a subset of E of size $r(M)$. Consider the square submatrix $A(X)[S]$, consisting of the columns of $A(X)$ indexed by S . If S is a basis of M , then $\det(A(X)[S])$

is not in $\langle Q \rangle$ and has degree at most n by the choice of n . Thus $\det(A(f(X))[S]) = \det(A(\alpha_1, \dots, \alpha_{|X|})[S])$, the evaluation of this polynomial at $\alpha_1, \dots, \alpha_{|X|} \in \mathbb{F}$, is non-zero by choice of $\alpha_1, \dots, \alpha_{|X|}$. Conversely, if S is dependent in M , then $\det(A(X)[S])$ is in $\langle Q \rangle$, and by the choice of $\alpha_1, \dots, \alpha_{|X|} \in \mathbb{F}$, we have $\det(A(f(X))[S]) = \det(A(\alpha_1, \dots, \alpha_{|X|})[S]) = 0$. Thus the evaluation $A(\alpha_1, \dots, \alpha_{|X|})$ is a \mathbb{F} -representation of M . \square

We may thus evaluate the indeterminates in a representation $A(X)$ at elements from a field \mathbb{F} that respect the relations in Q and are otherwise sufficiently algebraically independent.

2.0.3 Encoding values as points on lines

We first look at how we will encode values of a field. For simplicity, this will be done as “points” on a “line” in the matroid. For a represented matroid, the number of vectors required to specify a flat is the rank of that flat. So for a matroid, a *point* is a rank-1 flat, while a *line* is a rank-2 flat. A matroid is *simple* when every point only consists of a single element.

We will think of a labelling of a simple line L using a field \mathbb{F} as an injective map from L to $\mathbb{F} \cup \{\infty\}$ or $\mathbb{F} \cup \{-\infty\}$. When an element α of $\mathbb{F} \setminus \{-\infty, \infty\}$ is being used as a label, we will denote it $[\alpha]$ to avoid confusion. We often only care about an element because of the algebraic connotation provided by a labelling. To simplify notation, we will often refer to an element using its label on a given line.

Let $[0]$ and $[\infty]$ label two fixed non-parallel elements on a line L in a matroid M . Note that $\{[0], [\infty]\}$ will form a circuit with each element on L that is not parallel to either $[0]$ or $[\infty]$. For a fixed representation f of M , we will label an element on L by $[\alpha]$ when $f([\alpha]) = \lambda_\alpha (f([0]) + \alpha f([\infty]))$, for some non-zero scalar λ_α , see Figure 2.1. We say that the point with elements labelled $[\alpha]$ is the value α *encoded on the line L with respect to $[0]$ and $[\infty]$* . We may instead encode with respect to $[0]$ and a non-parallel element $[-\infty]$ where an element on L is labelled by $[\alpha]$ when $f([\alpha]) = \lambda_\alpha (f([0]) - \alpha f([-\infty]))$ for some non-zero scalar λ_α . This is the correspondence used for representations of gain graphs and Dowling geometries, see [34, 6.10].

Note we may swap between the labelling $[\infty]$ and $[-\infty]$ by negating the corresponding column. Also note that row equivalent representations have the same encoding for each point. However, when considering projectively equivalent representations we may scale the representations of $[0]$ and $[\pm\infty]$ and thus scale all other encodings on L by a common factor.

$$\begin{array}{ccc} [0] & [\infty] & [\alpha] \\ \left(\begin{array}{ccc} 1 & 0 & 1 \\ 0 & 1 & \alpha \end{array} \right). \end{array}$$

Figure 2.1: A representation restricted to possible elements of L , up to row equivalence.

We may simultaneously encode values on multiple lines in a matroid. We say that encodings that come from the same representation are *kindred*. At times a labelling will give us sufficient information to obtain a representation of the matroid that encodes this labelling, if it exists. For a field \mathbb{F} , we say that a labelling is \mathbb{F} -*consistent* or *consistent* when there exists an \mathbb{F} -representation that encodes this labelling.

2.0.4 Algebraic relations through restrictions

We will encode algebraic relations by “connecting” geometric gadgets to lines that encode values. Imposing algebraic structure can be done with two flavours of gadgets, local and global. Local algebraic structures encode a single relation between values encoded on the same line, while global algebraic structures encode many similar relations between values in kindred encodings on a collection of lines.

2.1 Local algebraic structure; von Staudt constructions

Here we look at matroids that encode a single algebraic relation on a single line. We would like to use gadgets that impose a relation on the representations of certain points but do not otherwise interact with the line. For a matroid M , we say that a subset D of $E(M)$ is *modular*, when $r(D) + r(F) = r(D \cup F) + r(D \cap F)$ for all flats F in M . Intuitively, a modular set D only interacts with other structure of the matroid at subsets of D .

We first see constructions to enforce nondegenerate additive and multiplicative relations. We then see the squaring or inversion relation as a degenerate case of the multiplicative relation. Similarly, the doubling and negation relations are degenerate cases of the additive relation. More complicated algebraic relations may require intermediate calculations and “glueing” multiple gadgets to a single line, see Section 2.3.

2.1.1 Non-degenerate addition

Let \mathcal{O}^+ be the matroid with representation

$$A(x, y) = \begin{pmatrix} [0] & [\infty] & [x] & [y] & [x+y] & i_x & i_y & o_x & o_y \\ 1 & 0 & 1 & 1 & 1 & 0 & 1 & 1 & 0 \\ 0 & 1 & x & y & x+y & 0 & 0 & x & -y \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & 1 \end{pmatrix}$$

over $\mathbb{Q}(x, y)$. It is not difficult to check that $L = \{[0], [\infty], [x], [y], [x+y]\}$ is a modular line in \mathcal{O}^+ .

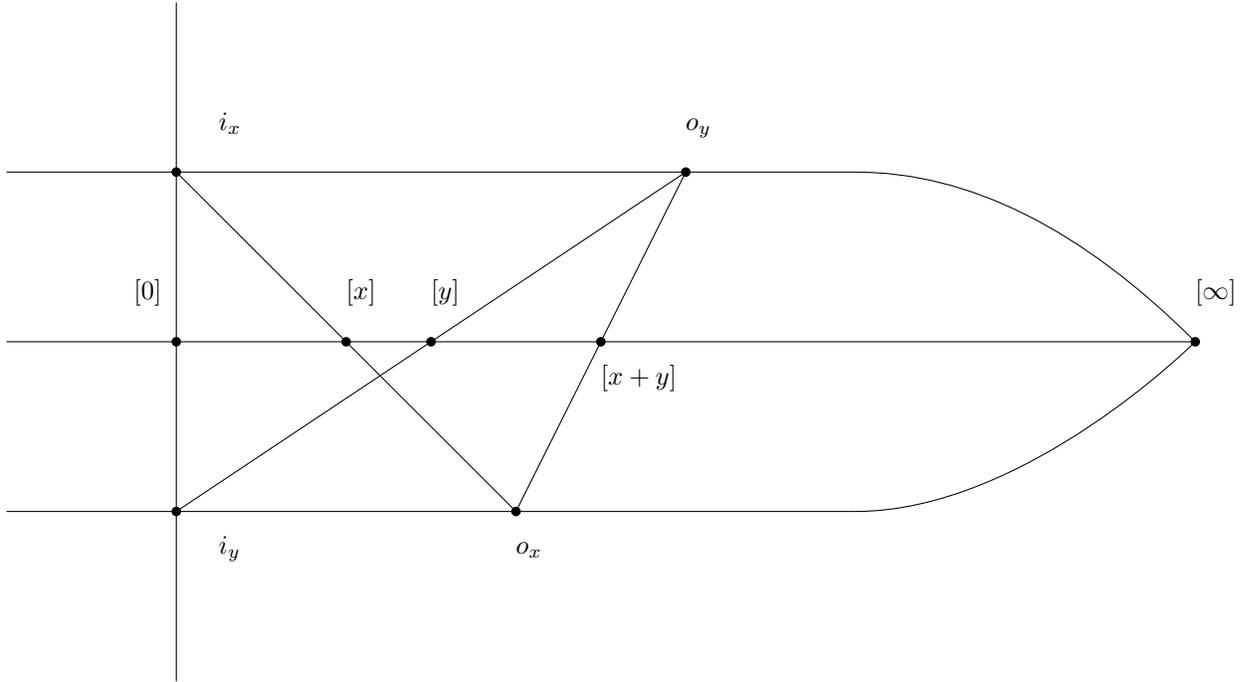


Figure 2.2: \mathcal{O}^+

Lemma 2.1.1. *Let $\bar{x}, \bar{y}, \bar{w}$ be distinct nonzero elements in some field \mathbb{F} . Then \mathcal{O}^+ has a representation A over \mathbb{F} such that*

$$A|L = \begin{pmatrix} [0] & [\infty] & [x] & [y] & [x+y] \\ 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & \bar{x} & \bar{y} & \bar{w} \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

if and only if $\bar{w} = \bar{x} + \bar{y}$.

Proof. If $\bar{w} = \bar{x} + \bar{y}$, then $A = A(\bar{x}, \bar{y})$ is the appropriate representation over \mathbb{F} . Conversely, suppose we have such a representation A . We will extend $A|L$ to uniquely determine A up to projective equivalence. As i_x is not contained in the flat L , without losing generality we may assume that $A|i_x = (0, 0, 1)^T$ through appropriate row operations. As $\{[0], i_x, i_y\}$ is a circuit, we may also assume that $A|i_y = (1, 0, 1)^T$ by scaling the third row and the columns corresponding to i_x and i_y . As o_x is spanned by $\{[x], i_x\}$ and by $\{[\infty], i_y\}$, we may assume $A|o_x = (1, \bar{x}, 1)^T$ through scaling. Similarly, as o_y is spanned by $\{[y], i_y\}$ and by $\{[\infty], i_x\}$, we may assume $A|o_y = (0, -\bar{y}, 1)^T$ through scaling. Finally, as $[x + y]$ is spanned by $\{o_x, o_y\}$ and by $\{[0], [\infty]\}$, we have that $(1, \bar{w}, 0)^T = A|[x + y] = (1, \bar{x} + \bar{y}, 0)^T$. Thus $\bar{w} = \bar{x} + \bar{y}$, as we wanted to show. \square

2.1.2 Non-degenerate multiplication

Let \mathcal{O}^* be the matroid with representation

$$P(x, y, z) = \begin{pmatrix} [0] & [\infty] & [x] & [y] & [z] & [xy/z] & p & o_{yz} & i_{xz} & o_{xw} & i_{yw} \\ 1 & 0 & 1 & 1 & 1 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & x & y & z & xy/z & 0 & 0 & -z & 0 & -y \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & x/z & 1 \end{pmatrix}$$

over $\mathbb{Q}(x, y)$. It is not difficult to check that $L = \{[0], [\infty], [x], [y], [z], [xy/z]\}$ is a modular line in \mathcal{O}^* .

Lemma 2.1.2. *Let $\bar{x}, \bar{y}, \bar{z}, \bar{w}$ be distinct nonzero elements in some field \mathbb{F} . Then \mathcal{O}^* has a representation P over \mathbb{F} such that*

$$P|L = \begin{pmatrix} [0] & [\infty] & [x] & [y] & [z] & [xy/z] \\ 1 & 0 & 1 & 1 & 1 & 1 \\ 0 & 1 & \bar{x} & \bar{y} & \bar{z} & \bar{w} \\ 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

if and only if $\bar{w}\bar{z} = \bar{x}\bar{y}$.

Proof. If $\bar{w}\bar{z} = \bar{x}\bar{y}$, then $P = P(\bar{x}, \bar{y}, \bar{z})$ is the appropriate representation over \mathbb{F} . Conversely, suppose we have such a representation P . We will extend $P|L$ to uniquely determine P up to projective equivalence. As p is not contained in the flat L , without

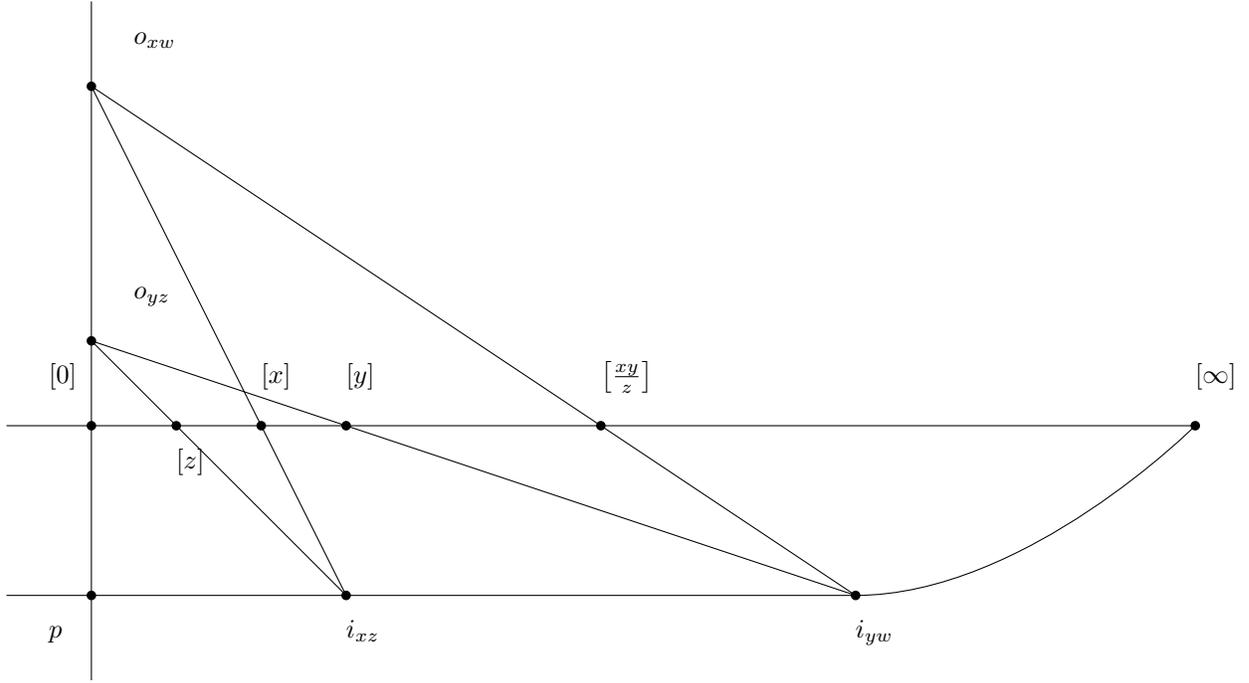


Figure 2.3: \mathcal{O}^*

losing generality we may assume that $P|p = (0, 0, 1)^T$ through appropriate row operations. As $\{[0], p, o_{yz}\}$ is a circuit, we may also assume that $P|o_{yz} = (1, 0, 1)^T$ by scaling the third row and the columns corresponding to p and o_{yz} . As i_{xz} is spanned by $\{[\infty], p\}$ and $\{[z], o_{yz}\}$, we may assume $P|i_{xz} = (0, -\bar{z}, 1)^T$ through scaling. Similarly, as i_{yw} is spanned by $\{[\infty], p\}$ and $\{[y], o_{yz}\}$, we may assume $P|i_{yw} = (0, -\bar{y}, 1)^T$ through scaling. And as o_{xw} is spanned by $\{[0], p\}$ and $\{[x], i_{xz}\}$, we may assume $P|o_{xw} = (1, 0, \bar{x}/\bar{z})^T$ through scaling. Finally, as $[xy/z]$ is spanned by $\{[0], [\infty]\}$ and $\{i_{yw}, o_{xw}\}$, we have that $(1, \bar{w}, 0)^T = P|[xy/z] = (1, \bar{x}\bar{y}/\bar{z}, 0)^T$. Thus $\bar{w}\bar{z} = \bar{x}\bar{y}$, as we wanted to show. \square

2.1.3 Squaring/Inversion

We want to be able to consider squaring and inversion relations. However, $[x]$ and $[y]$ lie on different points of \mathcal{O}^* which enforces the condition that $x \neq y$, even after evaluation in a specific field. We instead replace y with x in the matrix $P(x, y, z)$ from the previous section and get rid of the duplicated vector. Let \mathcal{O}' be the matroid represented by this

new matrix

$$Q(x, z) = \begin{pmatrix} [0] & [\infty] & [x] & [z] & [xx/z] & p & o_{xz} & i_{xz} & o_{xw} & i_{xw} \\ 1 & 0 & 1 & 1 & 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & x & z & xx/z & 0 & 0 & -z & 0 & -x \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 1 & x/z & 1 \end{pmatrix}$$

over $\mathbb{Q}(x, z)$. It is not difficult to check that $L = \{[0], [\infty], [x], [z], [xx/z]\}$ is a modular line \mathcal{O}' .

By taking the proof of Lemma 2.1.2 from the previous section, and replacing the indeterminate y with x and matrices P with Q , we get a proof of the following.

Lemma 2.1.3. *Let $\bar{x}, \bar{z}, \bar{w}$ be distinct nonzero elements in some field \mathbb{F} . Then \mathcal{O}' has a representation P over \mathbb{F} such that*

$$Q|L = \begin{pmatrix} [0] & [\infty] & [x] & [z] & [xx/z] \\ 1 & 0 & 1 & 1 & 1 \\ 0 & 1 & \bar{x} & \bar{z} & \bar{w} \\ 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

if and only if $\bar{w}\bar{z} = \bar{x}\bar{x}$.

Similarly, we can construct matroids \mathcal{O}^- and \mathcal{O}^2 that enforce the negation and doubling relations, respectively. A further degeneracy, to enforce the relation $x + x = 0$ with $x \neq 0$, yields the ‘‘Fano matroid’’, as depicted in Figure 1.1(ii).

2.1.4 Algebraic extensions

In Section 2.3, we will see that we can ‘‘glue’’ von Staudt matroids along their modular line to other matroids. By attaching copies of \mathcal{O}^+ , \mathcal{O}^* , \mathcal{O}' , \mathcal{O}^- , and \mathcal{O}^2 to a single line, we can enforce more complicated algebraic relations using intermediate calculations.

For example, consider the algebraic relation $x^2 + 1 = 0$. Let L be a line with non-parallel elements labelled $[0]$, $[1]$, $[x]$, $[xx] = [-1]$, and $[\infty]$ over the field $\mathbb{Q}[x]/\langle x^2 + 1 \rangle$. We can enforce the multiplicative relation between the labels $[1]$, $[x]$, and $[xx]$ by appropriately attaching \mathcal{O}' to this line. Similarly, the matroid \mathcal{O}' can be used to enforce the additive relation between the labels $[0]$, $[1]$, and $[-1]$. Let M be the resulting matroid ‘‘amalgam’’.

Consider a representation f of M over some field \mathbb{F} . By projective equivalence, we may assume that f correctly encodes $[1]$ on L with respect to $[0]$ and $[\infty]$. Together, the

gadgets \mathcal{O}' and \mathcal{O}'' we have used enforce that the element labelled $[x]$ encodes a solution to $x^2 + 1 = 0$ in \mathbb{F} . However, note that we have also inadvertently imposed the condition that $0, 1, x$, and $x^2 = -1$ are distinct, as they labelled non-parallel elements in L . Thus we have imposed the additional constraint that the characteristic of \mathbb{F} is not two, for instance. However, as a result of Lemma 2.3.4, the matroid M is representable in fields where $x^2 + 1$ is irreducible and has a root.

We have the following generalizations for fields of characteristic zero.

Theorem 2.1.4 (MacLane [24, Theorem 3]). *Let \mathbb{F} be a finite extension of the rationals. There exists a matroid $M_{\mathbb{F}}$ such that $M_{\mathbb{F}}$ is representable over an extension \mathbb{F}' of the rationals if and only if $\mathbb{F}' \supseteq \mathbb{F}$.*

One can check that there is an analogue of this result for each characteristic.

2.2 Global algebraic structure

Here we look at matroids that encode a set of algebraic relations on a collection of lines. Let $\mathcal{L} = \{L_1, \dots, L_n\}$ be a collection of lines with a point denoted $[\infty]_i$ or $[-\infty]_i$ on each line $L_i \in \mathcal{L}$. We say that $T = \{a_1, \dots, a_n\}$ is a *transversal* of \mathcal{L} when $a_i \in L_i - \{[\infty]_i, [-\infty]_i\}$ for each i . For certain kindred encodings of \mathcal{L} , the set \mathcal{T} of dependent transversals will encode which values satisfy a prescribed algebraic relation.

2.2.1 Sums; Spikes

Say that we would like to encode within a matroid a relation of the form $\sum_{i=1}^n \alpha_i = 0$. For a given set of values $\{\alpha_1, \dots, \alpha_n\}$, consider the matrix of the form

$$\left(\begin{array}{cccc|c} 1 & 0 & \dots & 0 & -1 \\ 0 & 1 & \dots & 0 & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -1 \\ \hline \alpha_1 & \alpha_2 & \dots & \alpha_{n-1} & \alpha_n \end{array} \right). \quad (2.1)$$

Note this matrix is singular precisely when $\sum_{i=1}^n \alpha_i = 0$.

To make use of this fact, we would like to encode the values α_i appropriately. We need to be able to determine each line L_i and the points that would correspond to $[0]_i$ and $[\infty]_i$

on L_i . Consider introducing vectors corresponding to the points $[0]_1, \dots, [0]_n$ as in the matrix below. Note that for each $i \in \{1, \dots, n\}$, if $\alpha_i \neq 0$ then the points $[0]_i$ and $[\alpha_i]_i$ span a line L_i in the represented matroid.

$$\left(\begin{array}{cccccc|cccccc} [0]_1 & [0]_2 & \dots & [0]_{n-1} & [0]_n & [\alpha_1]_1 & [\alpha_2]_2 & \dots & [\alpha_{n-1}]_{n-1} & [\alpha_n]_n \\ \hline 1 & 0 & \dots & 0 & -1 & 1 & 0 & \dots & 0 & -1 \\ 0 & 1 & \dots & 0 & -1 & 0 & 1 & \dots & 0 & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -1 & 0 & 0 & 0 & 1 & -1 \\ \hline 0 & 0 & \dots & 0 & 0 & \alpha_1 & \alpha_2 & \dots & \alpha_{n-1} & \alpha_n \end{array} \right)$$

Suppose we have a representation of a matroid that is projectively equivalent to the above. As long as we have at least two lines $L_i = \{[0]_i, [\alpha_i]_i\}$ and $L_j = \{[0]_j, [\alpha_j]_j\}$, we can determine the representation of a new element, p , that would lie in their intersection. With the representation above, p is represented by the n -th standard basis vector, \mathbf{e}_n . For each $i \in \{1, \dots, n\}$, we now have that the label $[\alpha_i]_i$ actually encodes α_i on the line L_i with $[0]_i$ and $[\infty]_i = p$. So, as we can retrieve the appropriate $[\infty]_i$ for each line, a matroid that has a representation projectively equivalent to the matrix above can be used to encode whether or not $\sum_{i=1}^n \alpha_i$ is zero.

We now see matroids that can each encode a family of relations of the form $\sum_{i=1}^n \alpha_i = 0$. However, such a matroid may be non-representable if the relations it would encode are inconsistent (see Theorem 5.3.1 for an example).

Spike-like matroids

A *spike-like matroid on n lines with tip p* is a simple matroid that is the union of a set of lines $\{L_1, \dots, L_n\}$ which all contain p and where, for any transversal $T = \{a_1, \dots, a_n\}$ with a_i in *leg* $L_i - \{p\}$, we have that each $(T \cup \{p\}) - \{a_i\}$ is a basis. Thus each transversal is either a basis or a circuit-hyperplane. We say that a matroid is *spike-like* on these n lines when it is a restriction of such a matroid. It is not hard to check that a spike-like matroid is determined by its lines L_1, \dots, L_n and its dependent transversals. Furthermore, for any set \mathcal{T} of transversals in which no two transversals differ on only a single leg, there is a spike-like matroid with \mathcal{T} as its set of dependent transversals, see [34, Proposition 1.5.17].

Representations

Let $\Lambda = (E, \mathcal{C})$ be a representable spike-like matroid with at least four legs, each containing at least two points. Let f be a representation of Λ over a field \mathbb{F} . As we have four legs that

define lines of Λ , we can determine a representation of a tip p for Λ . Let $T = \{a_1, \dots, a_n\}$ be a dependent transversal and fix an element of T , say a_n for convenience. As $(T \cup \{p\}) - \{a_n\}$ is a basis, by row operations and nonzero column scaling we may assume that this basis is mapped to the standard basis vectors. For $i \in \{1, \dots, n-1\}$, let a_i be represented by the i -th standard basis vector \mathbf{e}_i and label a_i by $[0]_i$. Let p be represented by the n -th standard basis vector \mathbf{e}_n . Note that a_n is represented by $\sum_{i=1}^{n-1} \mathbf{e}_i$, denoted $\mathbb{K} - \mathbf{e}_n$ and label a_n by $[0]_n$. By column scaling, we may assume that, for $i \in \{1, \dots, n-1\}$, each element $x_i \in L_i - \{p\}$ is represented by $\mathbf{e}_i + \alpha_i \mathbf{e}_n$ for some $\alpha_i \in \mathbb{F}$, and we henceforth label x_i as $[\alpha_i]_i$. Similarly, each $x_n \in L_n - \{p\}$ is represented by $(\mathbb{K} - \mathbf{e}_n) + \alpha_n \mathbf{e}_n$ for some $\alpha_n \in \mathbb{F}$, and we henceforth label x_n as $[\alpha_n]_n$. When this is done, we will say this encoding is *with respect to the dependent transversal T* . Recalling (2.1), we have that a transversal $T' = \{[\alpha_1]_1, \dots, [\alpha_n]_n\}$ is dependent if and only if $\sum_{i=1}^n \alpha_i = 0$.

There is some freedom in scaling when converting to a representation of this form. Specifically, we may scale all values α_i of the encoding by a common non-zero factor $\lambda \in \mathbb{F} - \{0\}$. Up to this scaling factor, each representation of Λ is projectively equivalent to a unique encoding with respect to T .

2.2.2 Products; Swirls

Say that we would like to encode within a matroid relations of the form $\prod_{i=1}^n \alpha_i = 1$. For a given set of values $\{\alpha_1, \dots, \alpha_n\}$, consider the matrix of the form

$$\begin{pmatrix} 1 & 0 & 0 & \dots & 0 & -\alpha_n \\ -\alpha_1 & 1 & 0 & \dots & 0 & 0 \\ 0 & -\alpha_2 & 1 & \ddots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & 1 & 0 \\ 0 & 0 & 0 & \dots & -\alpha_{n-1} & 1 \end{pmatrix}. \quad (2.2)$$

Note that this matrix has determinant $1 - \prod_{i=1}^n \alpha_i$. Thus the above matrix is singular precisely when $\prod_{i=1}^n \alpha_i = 1$.

To make use of this fact, we would like to encode the values α_i appropriately. We need to be able to determine each line L_i and the points that would correspond to $[0]_i$ and $[-\infty]_i$ on L_i . Consider introducing vectors corresponding to the points $[1]_1, \dots, [1]_n$ as in the matrix below. Note that for each $i \in \{1, \dots, n\}$, if $\alpha_i \neq 1$ then the points $[1]_i$ and $[\alpha_i]_i$

span a line L_i in the represented matroid.

$$\begin{pmatrix} [1]_1 & [1]_2 & \cdots & [1]_{n-1} & [1]_n & [\alpha_1]_1 & [\alpha_2]_2 & \cdots & [\alpha_{n-1}]_{n-1} & [\alpha_n]_n \\ \left(\begin{array}{cccc|cccccc} 1 & 0 & \cdots & 0 & -1 & 1 & 0 & \cdots & 0 & -\alpha_n \\ -1 & 1 & \ddots & \vdots & 0 & -\alpha_1 & 1 & \ddots & \vdots & 0 \\ 0 & -1 & \ddots & 0 & \vdots & 0 & -\alpha_2 & \ddots & 0 & \vdots \\ \vdots & 0 & \ddots & 1 & 0 & \vdots & 0 & \ddots & 1 & 0 \\ 0 & \vdots & \cdots & -1 & 1 & 0 & \vdots & \cdots & -\alpha_{n-1} & 1 \end{array} \right) \end{pmatrix}.$$

Suppose we have a representation of a matroid that is projectively equivalent to the above. If we have two consecutive lines $L_{i-1} = \{[1]_{i-1}, [\alpha_{i-1}]_{i-1}\}$ and $L_i = \{[1]_i, [\alpha_i]_i\}$, we determine the representation of a new element, b_i , that lies in the intersection of L_{i-1} and L_i . With the representation above, b_i is represented by the i -th standard basis vector, \mathbf{e}_i . For each $i \in \{1, \dots, n\}$, we now have that the label $[\alpha_i]_i$ actually encodes α_i on the line L_i with $[0]_i = b_i$ and $[-\infty]_i = b_{i+1}$. So, as we can retrieve the appropriate $[0]_i$ and $[-\infty]_i$ for each line, a matroid that has a representation projectively equivalent to the matrix above can be used to encode whether or not $\prod_{i=1}^n \alpha_i$ is one.

We now see matroids that can each encode a family of relations of the form $\prod_{i=1}^n \alpha_i = 1$. However, such a matroids may be non-representable if the relations it would encode are inconsistent.

Swirl-like matroids

A *swirl-like matroid on n lines with joints b_1, \dots, b_n* is a simple matroid that is the union of a set of lines $\{L_i | i \in \mathbb{Z}_n\}$ where $\{b_i | i \in \mathbb{Z}_n\}$ is a basis and each line L_i contains b_i and b_{i+1} . Note that each transversal $T = \{a_1, \dots, a_n\}$ with a_i in *edge* $L_i - \{b_i, b_{i+1}\}$ is either a basis or a circuit-hyperplane. We say that a matroid is *swirl-like* on these n lines when it is a restriction of such a matroid. It is not hard to check that a swirl-like matroid is determined by its lines L_1, \dots, L_n and its dependent transversals. Furthermore, for any set \mathcal{T} of transversals in which no two transversals differ on only a edge, there is a swirl-like matroid with \mathcal{T} as its set of dependent transversals.

Representations

Let $\Omega = (E, \mathcal{C})$ be a representable swirl-like matroid with joints b_1, \dots, b_n . Let f be a representation of Ω over a field \mathbb{F} . As $B = \{b_i | i \in \mathbb{Z}_n\}$ is a basis, by row operations

and column scaling, we may assume that, for each $i \in \mathbb{Z}_n$, the joint b_i is represented by the i -th standard basis vector \mathbf{e}_i . By column scaling, we may assume that each element $a_i \in L_i - \{b_i, b_{i+1}\}$ is represented by $\mathbf{e}_i - \alpha_i \mathbf{e}_{i+1}$ for some $\alpha_i \in \mathbb{F} - \{0\}$, and we henceforth label a_i as $[\alpha_i]_i$. Recalling (2.2), we have that a transversal $T = \{[\alpha_1]_1, \dots, [\alpha_n]_n\}$ is dependent if and only if $\prod_{i=1}^n \alpha_i = 1$.

There is some freedom in scaling when converting to a representation of this form. For a fixed dependent transversal $T = \{[\alpha_1]_1, \dots, [\alpha_n]_n\}$ of Ω , we can scale the representation so that $\alpha_i = 1$ for each $i \in \mathbb{Z}_n$. When this is done, we will say this encoding is *with respect to the dependent transversal T* . Each representation of Ω is projectively equivalent to a unique encoding with respect to a given dependent transversal T .

2.3 Sharing structure; Amalgams

The matroid structures from the previous sections each have limited use algebraically. However, we will now see how we can “glue” matroids together.

Say we have two matroids M_1 and M_2 on ground sets E_1 and E_2 , respectively. An *amalgam* of M_1 and M_2 is a matroid M with ground set $E_1 \cup E_2$ such that $M|_{E_1} = M_1$ and $M|_{E_2} = M_2$.

For this to exist, M_1 and M_2 must have compatible structure over the set $L = E_1 \cap E_2$ which the amalgam occurs *across*. An immediate necessary condition, for instance, is that M_1 and M_2 have the same structure at L , that is, a common restriction $R = M_1|_L = M_2|_L$. However, this may not be sufficient.

Recall that for a matroid M , a subset L of $E(M)$ is *modular*, when

$$r(L) + r(F) - r(L \cup F) = r(L \cap F) \tag{2.3}$$

for all flats F in M . For two sets X and Y of a matroid M , the value $r(X) + r(Y) - r(X \cup Y)$, is the *local connectivity* between X and Y , denoted $\square(X, Y)$. Intuitively, this would be the rank of the intersection of $\text{cl}(X)$ and $\text{cl}(Y)$ if M was actually a vector space, see Figure 2.4. Noting $F = \text{cl}(F)$ for each flat F in (2.3), we can interpret a modular set, L , as one which behaves as a vector subspace in how it intersects flats of the matroid.

Brylawski [6], showed that if $L = E_1 \cap E_2$ is modular in one of M_1 or M_2 then we can define an amalgam. Specifically, if $M_1 = (E_1, \mathcal{F}_1)$ and $M_2 = (E_2, \mathcal{F}_2)$ are matroids described by their flats, and $R = M_1|_L = M_2|_L$ with $L = E_1 \cap E_2$ modular in one of M_1 or M_2 , then

$$\mathcal{F} = \{F \subseteq E_1 \cup E_2 : F \cap E_1 \in \mathcal{F}_1 \text{ and } F \cap E_2 \in \mathcal{F}_2\} \tag{2.4}$$

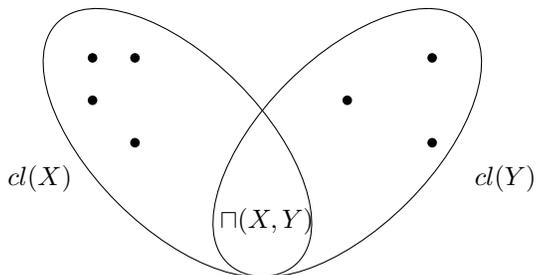


Figure 2.4: Local connectivity in a representable matroid

is the set of flats of an amalgam of M_1 and M_2 , see [34, Proposition 11.4.13]. This amalgam is known as the *general parallel connection* of M_1 and M_2 , denoted $P_R(M_1, M_2)$. If $L = E_1 \cap E_2 = \emptyset$, then we call the general parallel connection of M_1 and M_2 the *direct sum* of M_1 and M_2 , denoted $M_1 \oplus M_2$.

The general parallel connection is an example of a “proper” amalgam. An amalgam M of M_1 and M_2 is the *proper* amalgam when, for every flat F of M , the rank function r of M satisfies

$$r(F) = r(F \cap E_1) + r(F \cap E_2) - r(F \cap L), \quad (2.5)$$

see [34, Theorem 11.4.3]. In other words, $r(F \cap L) = \sqcap(F \cap E_1, F \cap E_2)$, so intuitively flats can only cross between E_1 and E_2 within $L = E_1 \cap E_2$. More generally, when the proper amalgam of $M_1 = (E_1, r_1)$ and $M_2 = (E_2, r_2)$ exists, it is denoted $M_1 \oplus_R M_2$ where $R = M_1|L = M_2|L$ and has rank function given by

$$r(X) = \min\{r_1(S \cap E_1) + r_2(S \cap E_2) - r_1(S \cap L) : X \subseteq S \subseteq E_1 \cup E_2\} \quad (2.6)$$

for a subset X of $E_1 \cup E_2$ [34, Theorem 11.4.2]. Besides the instance of general parallel connections, the proper amalgam also exists when all the flats of $R = M_1|L = M_2|L$ are modular [34, Theorem 11.4.10]. This is the case when R is a line, but more generally R can be the direct sum of finite projective geometries, see [4, pp 90–93]. When the notation $M_1 \oplus_R M_2$ is used, we will assume that elements of $E(M_1) - E(R)$ have been renamed if necessary to ensure that $E(M_1) \cap E(M_2) = E(R)$.

Recall that a hyperplane of a matroid $M = (E, r)$ is a flat H for which $r(H) = r(M) - 1$. By (2.4) and (2.5) and the observation that $r(P_R(M_1, M_2)) = r(M_1) + r(M_2) - r(D)$, we have the following.

Remark 2.3.1. *Let $M_1 = (E_1, \mathcal{F}_1)$ and $M_2 = (E_2, \mathcal{F}_2)$ be matroids with $R = M_1|L = M_2|L$ and $L = E_1 \cap E_2$ modular in one of M_1 or M_2 . Then H is a hyperplane of $P_R(M_1, M_2)$ if and only if*

- $H \supseteq E_1$ and $H \cap E_2 \supset L$ is a hyperplane of M_2 ,
- $H \supseteq E_2$ and $H \cap E_1 \subset L$ is a hyperplane in M_1 , or
- $H \not\supseteq L$ and $H \cap E_1$ and $H \cap E_2$ are hyperplanes in M_1 and M_2 , respectively.

Let $M = (E, r)$ be a representable matroid with matrix representation A , that is, $r(X) = \text{rank}(A|X)$ for $X \subseteq E$. Let H be a hyperplane of M , that is, a maximal set of rank $r(M) - 1$. Thus $\text{rank}(A|H) = r(H) = r(X) - 1 = \text{rank}(A|X) - 1 = \text{rank}(A) - 1$ for any $X \supseteq M$ that properly contains H . By the rank-nullity theorem, there is a vector \mathbf{w}_H in the nullspace of $(A|H)^T$ that is orthogonal to precisely the columns that are indexed by elements of H . We can summarize this as the follows.

Remark 2.3.2. *If M is a representable matroid with representation $\{\mathbf{v}_e\}_{e \in E(M)}$, then for any hyperplane H of M there is a vector \mathbf{w}_H such that $H = \{e \in E(M) : (\mathbf{w}_H)^T \mathbf{v}_e = 0\}$.*

By using the hyperplane description of general parallel connection, Remark 2.3.1, in conjunction with the characterizing vectors of hyperplanes, Remark 2.3.2, we get the following.

Theorem 2.3.3. *Let \mathbb{F} be a field. Let M_1 and M_2 be matroids on ground sets E_1 and E_2 , respectively. Let $L = E_1 \cap E_2$ be a common modular line in M_1 and M_2 with $R = M_1|L = M_2|L$. The general parallel connection $P_R(M_1, M_2)$ has \mathbb{F} -representation A with $A|E_1$ and $A|E_2$ row equivalent to A_1 and A_2 , respectively, if and only if M_1 and M_2 have \mathbb{F} -representations A_1 and A_2 , respectively, with $A_1|L$ and $A_2|L$ row equivalent to each other.*

By first using “principal extensions” (see Section 4.0.2) to ensure two lines have the same elements, we have the following more general result.

Lemma 2.3.4. *Let \mathbb{F} be a field. Let M_1 and M_2 be simple matroids with \mathbb{F} -representations A_1 and A_2 and modular lines L_1 and L_2 , respectively. Let $L = L_1 \cap L_2 = E(M_1) \cap E(M_2)$ and have common restriction $R = M_1|L = M_2|L$ of rank 2. Let A' be a matrix over \mathbb{F} that represents a simple rank-2 matroid on ground set $L_1 \cup L_2$. If $A'|L_1$ is row equivalent to $A_1|L_1$ and $A'|L_2$ is row equivalent to $A_2|L_2$ then there is a \mathbb{F} -representation A of $M_1 \oplus_R M_2$ for which $A|E(M_1) = A_1$ and $A|E(M_2) = A_2$.*

In other words, if M_1 and M_2 are matroids that have consistent \mathbb{F} -encodings on modular lines L_1 and L_2 , respectively, such that the elements in $L = L_1 \cap L_2$ have the same labelling in L_1 and L_2 with the only labels that occur for both, then there is a consistent \mathbb{F} -encoding of the line $L_1 \cup L_2$ in $M_1 \oplus_R M_2$ with this labelling. This lemma is what allows us to “glue” consistent gadgets together.

Proof of Lemma 2.3.4. We see that we may extend L_1 by elements in $L_2 - L_1$ in such a way that the only flats of M_1 that span these new elements are those that contain L_1 . Specifically, we represent each new element on the line L_1 according to the row equivalence of $A_1|L_1$ with the rank-2 matrix A' . As L_1 is modular in M_1 and A' represents a simple matroid, each new element only lies in flats spanning L_1 , so this extended line $L' = L_1 \cup L_2$ is modular in our new matroid M'_1 . Similarly, we can extend M_2 to a matroid M'_2 where L' is a modular line. Note A' represents $M'_1|L' = M'_2|L'$, which we denote R' . We may now consider the general parallel connection $P_{R'}(M'_1, M'_2)$ with ground set $E(M_1) \cup E(M_2)$. By construction, both $A'_1|L'$ and $A'_2|L'$ are row equivalent to A' . So by Theorem 2.3.3, $P_{R'}(M'_1, M'_2)$ is \mathbb{F} -representable. Note that $M_1 \oplus_R M_2$ and $P_{R'}(M'_1, M'_2)$ have the same ground set and the same rank function as they are both proper amalgams. Thus we have that $M_1 \oplus_R M_2 = P_{R'}(M'_1, M'_2)$ is \mathbb{F} -representable. □

2.3.1 Cannibalizing structure; Pinned extensions

Once we have shared structure by an amalgam across a line, there may be elements on the line that are “pointed at” from both sides of the amalgam, or “pinned”. We will see that if we remove an element that is “pinned” on the line, we can “reinsert” it to retrieve the original matroid. This will allow us to take advantage of useful structure in a matroid to get new constructions.

Recall that the *local connectivity* between two sets S and T of a matroid M is

$$\square_M(S, T) = r_M(S) + r_M(T) - r_M(S \cup T).$$

Now, for disjoint sets X, Y , and C in M , we have

$$\square_{M/C}(X, Y) = \square_M(X, Y \cup C) - \square_M(X, C),$$

which can be easily verified by expanding both sides. We will also use the fact that, if $\square_M(X, Y) = 0$ and e is spanned by both X and Y , then e is a loop; this follows since

$$\begin{aligned} r_M(\{e\}) &\leq r_M(\text{cl}_M(X) \cap \text{cl}_M(Y)) \\ &\leq r_M(\text{cl}_M(X)) + r_M(\text{cl}_M(Y)) - r_M(\text{cl}_M(X) \cup \text{cl}_M(Y)) \\ &= r_M(X) + r_M(Y) - r_M(X \cup Y) \\ &= \square_M(X, Y) = 0. \end{aligned}$$

Let (S_1, S_2) be a partition of the ground set of a matroid M' such that $\square(S_1, S_2) = 2$. This is called a *3-separation of M'* . Let M be obtained from M' by extending by a non-loop element e into the closures of both S_1 and S_2 . Unlike the case with 2-separations, this does not uniquely determine M . However, under some additional hypotheses, the following result shows that we can uniquely determine M .

Lemma 2.3.5. *Let e be a non-loop element of a matroid M , let (S_1, S_2) be a 3-separation of $M \setminus e$, and let $Y_1 \subseteq S_1$ and $Y_2 \subseteq S_2$ such that $\square_M(Y_1, S_2) = 1$, $\square_M(S_1, Y_2) = 1$, and e is spanned by both Y_1 and Y_2 in M . Then a flat F of M spans e if and only if either*

$$(i) \quad \square_M(F \cap S_1, Y_2) = 1 \text{ or } \square_M(Y_1, F \cap S_2) = 1, \text{ or}$$

$$(ii) \quad \square_M(F \cap S_1, S_2) = \square_M(S_1, F \cap S_2) = 1 \text{ and } \square_M(F \cap S_1, F \cap S_2) = 0.$$

Proof. Let $F_1 = F \cap S_1$ and $F_2 = F \cap S_2$. First, suppose that $\square_M(F_1, Y_2) = 1$. Then $\square_{M/F_1}(S_1 - F_1, Y_2) = \square_M(S_1, Y_2) - \square_M(F_1, Y_2) = 0$. However, e is in the closure of both $S_1 - F_1$ and Y_2 in M/F_1 . Thus e is a loop in M/F_1 and hence e is spanned by F . By symmetry, if $\square_M(Y_1, F_2) = 1$, then e is spanned by F .

Now suppose that $\square_M(F_1, S_2) = \square_M(S_1, F_2) = 1$ and $\square_M(F_1, F_2) = 0$. Then

$$\begin{aligned} \square_{M/F}(S_1 - F_1, S_2 - F_2) &= \square_{M/F_1}(S_1 - F_1, S_2) - \square_{M/F_1}(S_1 - F_1, F_2) \\ &= \square_M(S_1, S_2) - \square_M(F_1, S_2) \\ &\quad - \square_M(S_1, F_2) + \square_M(F_1, F_2) \\ &= 0. \end{aligned}$$

However, e is spanned by both $S_1 - F_1$ and $S_2 - F_2$ in M/F . Thus e is a loop in M/F and, hence, F spans e .

Conversely, suppose that F spans e and hence that e is a loop in M/F . We may assume that e is not spanned by either F_1 or F_2 since otherwise (i) holds. Since e is spanned by F_2 in M/F_1 , we have $\square_M(F_1, S_2) = 1$. Similarly $\square_M(S_1, F_2) = 1$. Moreover, again since e is spanned by F_2 in M/F_1 , we have $1 = \square_{M/F_1}(S_1 - F_1, F_2) = \square_M(S_1, F_2) - \square_M(F_1, F_2) = 1 - \square_M(F_1, F_2)$ and, hence $\square_M(F_1, F_2) = 0$, so (ii) holds. \square

When the hypotheses from the previous theorem are satisfied, we say that (Y_1, Y_2) *pins* e and that M is a *pinned extension into a 3-separation of $M \setminus e$* . Oxley characterized when such an extension exists in general [35], but the following is an easier case.

Lemma 2.3.6. *Let \mathbb{F} be a field. The class of \mathbb{F} -representable matroids is closed under pinned extensions into 3-separations.*

Proof. Let M be a matroid with \mathbb{F} -representation g . Let Y_1 and Y_2 be subsets of $E(M)$ with $\square(Y_1, Y_2) = 1$. By the modularity of the dimension of subspaces, the subspace spanned by $g(Y_1)$ intersects the subspace spanned by $g(Y_2)$. A non-trivial vector v in this intersection subspace gives us a representation of an extension point e that is spanned by both Y_1 and Y_2 . \square

Once we have a matroid N that is not real-representable, this lemma provides a technique to construct a new matroid that are not real-representable. Specifically, consider $M \setminus p$ where M is the amalgam of N with another matroid such that an element p of N is pinned across a 3-separation of M . The matroid $M \setminus p$ is not real-representable as otherwise M is real-representable by Lemma 2.3.6 and this contradicts the assumption that the restriction N is not real-representable.

Chapter 3

Axiomatization

This chapter is mostly [27] by Mayhew, Newman, and Whittle.

Here we consider using logical conditions to characterize real-representability. With a strong enough logical language, we can define whatever we can conceive of. Instead, we will use a natural independent-set language MS_0 . Properties that are finitely axiomatizable in this language are intuitively not too complicated from a combinatorial perspective. We will see that we can state matroid axioms with this language — as opposed to Vámos’s language [43] — and explore some matroid properties that we can also finitely axiomatize. Surprisingly, we will see that the prototypical class of matroids — real-representable matroids — is not so easily defined [27], even within the class of representable matroids.

3.0.1 The independent-set language MS_0

As matroids are combinatorial in nature, we will use a *second-order language*: one with variables for elements, but also for sets of tuples of elements. We use the independent-set formulation of matroids, as this requires only element and set variables — a *monadic* second-order language. We will use the language MS_0 given by Mayhew, Newman, and Whittle in [27], where, instead of element variables, there is a relation identifying singleton sets. This is equivalent to the language used by Hliněný in [15].

The language MS_0 will be used to finitely encode conditions on pairs (E, \mathcal{I}) , where E is a set of elements and \mathcal{I} is the set of “independent” subsets of E . The language MS_0 consists of: countably many variables X_1, X_2, \dots ; unary predicates Sing and Ind , and the binary predicate \subseteq ; the connectives \neg , \wedge , \vee , and \rightarrow ; and the quantifiers \exists and \forall . For

a pair (E, \mathcal{I}) , if the variable X_i is assigned to the subset A_i of E , then $\text{Sing}(X_i)$ is true precisely when A_i is a singleton, and $\text{Ind}(X_i)$ is true precisely when $A_i \in \mathcal{I}$.

3.1 Matroid axiomatization

We now recall the prototypical example of linear independence. Let \mathcal{I} be a collection of subsets of a finite set E . We have that (E, \mathcal{I}) is *representable* over a field \mathbb{F} when there is a multiset $A = \{v_e\}_{e \in E}$ of vectors from a vector space over \mathbb{F} such that \mathcal{I} corresponds to linearly independent subsets of A . Can we give criteria for when (E, \mathcal{I}) is representable? There are some fundamental conditions which \mathcal{I} must satisfy. For instance:

(I1): \mathcal{I} is not empty.

(I2): Any subset of an element of \mathcal{I} is also in \mathcal{I} .

(I3): If B is maximal in \mathcal{I} and S is in \mathcal{I} but not maximal, then there is an $e \in B - S$ for which there is T in \mathcal{I} that contains exactly the elements $t \in S \cup \{e\}$.

The condition (I3) is not a commonly used independence axiom and is closer to the “basis exchange” axiom, but we wish to avoid using the size of sets. Since this is done, we can encode each of these properties over the language MS_0 as a *sentence*, that is, a finite expression where all variables are quantified [27].

$$\begin{aligned} \phi_{(I1)} &: \exists X_1 \text{Ind}(X_1) \\ \phi_{(I2)} &: \forall X_2 \forall X_3 (\text{Ind}(X_2) \wedge (X_3 \subseteq X_2)) \rightarrow \text{Ind}(X_3) \\ \phi_{(I3)} &: \forall B \forall S (\text{Ind}(B) \wedge (\forall B_+ (B \subseteq B_+ \rightarrow (B_+ \subseteq B \vee \neg \text{Ind}(B_+)))) \wedge \\ &\quad \text{Ind}(S) \wedge (\exists S_+ \text{Ind}(S_+) \wedge (S \subseteq S_+) \wedge \neg(S_+ \subseteq S)) \rightarrow \\ &\quad (\exists \iota_e \text{Sing}(\iota_e) \wedge (\iota_e \subseteq B) \wedge \neg(\iota_e \subseteq S) \wedge \\ &\quad \exists T \text{Ind}(T) \wedge (\forall \iota_t \text{Sing}(\iota_t) \rightarrow (\iota_t \subseteq T \leftrightarrow ((\iota_t \subseteq S) \vee (\iota_t \subseteq \iota_e)))) \end{aligned}$$

For consistency with (I3) above, $X_4, X_5, X_6, X_7, X_8, X_9, X_{10}$ have been renamed in $\phi_{(I3)}$ as $B, S, B_+, S_+, \iota_e, T, \iota_t$ respectively.

Whitney named the class of structures (E, \mathcal{I}) that satisfy these conditions *matroids* [45]. Thus, the class of matroids is finitely axiomatizable over MS_0 . Equivalently, by considering the conjunction of the axioms, we have:

Remark 3.1.1 (Mayhew, Newman, Whittle [27]). *A pair (E, \mathcal{I}) is a matroid precisely when it satisfies the sentence $\phi_{(I1)} \wedge \phi_{(I2)} \wedge \phi_{(I3)}$.*

By definition, being a matroid is a necessary condition for a pair (E, \mathcal{I}) to be representable. We will see that it is impossible to expand this to a finite list of conditions that characterize representability. Equivalently, as we could take the conjunction of such a list:

[Theorem 3.3.1 (Mayhew, Newman, Whittle [27]). *There is no sentence ϕ in MS_0 such that a matroid is representable precisely when it satisfies ϕ .*

This is also the case when we restrict to real-representability.

Theorem 3.1.2 (Mayhew, Newman, Whittle [27]). *Let \mathbb{F} be an infinite field. There is no sentence $\phi_{\mathbb{F}}$ in MS_0 such that a matroid is \mathbb{F} -representable precisely when it satisfies $\phi_{\mathbb{F}}$.*

This was also shown to be the case for any infinite field, see Theorem 3.4.1. Even assuming that the matroids in question are representable does not necessarily improve the situation. In this chapter we prove the following strengthening of the real case of the previous result.

[Theorem 3.4.2 (Campbell)]. *There is no sentence ϕ in MS_0 such that a complex-representable matroid is real-representable precisely when it satisfies ϕ .*

That is to say, within the class of complex-representable matroids, real-representability is not finitely axiomatizable.

Before we get to the proofs of these theorems, we will see some matroid properties that can be defined with this independent-set language.

3.1.1 Important sets in a matroid

We can use MS_0 to define some useful types of sets in a matroid.

Bases

A subset S of E is a *basis* when it is maximally independent, that is, S is a basis when it satisfies the formula

$$\text{Ind}(S) \wedge (\forall X (S \subseteq X \rightarrow (X \subseteq S \vee \neg \text{Ind}(X))))$$

Circuits

A subset S of E is a *circuit* when it is minimally dependent, that is, S is a circuit when it satisfies the formula

$$(\neg \text{Ind}(S)) \wedge (\forall X (X \subseteq S \rightarrow (S \subseteq X \vee \text{Ind}(X))))$$

3.1.2 Minors

Given a set system (E, \mathcal{I}) , minor operations are the natural way to restrict the scope of the structure. Given an element e in E , we get a new set system $(E - \{e\}, \mathcal{I}')$ where \mathcal{I}' is all the sets in \mathcal{I} that do not contain e (*deletion*). Alternatively, we can “project” from e when $\{e\} \in \mathcal{I}$ by instead taking \mathcal{I}' to be all the sets $I - e$ where $e \in I \in \mathcal{I}$ (*contraction of e* , denoted M/e). When there is no set in \mathcal{I} containing e , we define $M/e = M \setminus e$. When a matroid N can be obtained from M with a sequence of these operations, we call N a *minor* of M .

For a fixed matroid N on ground set $\{1, \dots, n\}$ and with \mathcal{I}' as its set of independent sets, a matroid M contains a minor isomorphic to N precisely when M satisfies:

There exist disjoint singleton sets X_1, \dots, X_n and a disjoint independent set X_{n+1} (the set to be contracted) such that:

- for each $\{i_1, \dots, i_t\} \in \mathcal{I}'$, there is an independent set Y for which a singleton is a subset of Y precisely when this singleton is a subset of one of $X_{i_1}, \dots, X_{i_t}, X_{n+1}$ (namely, Y is the set $X_{i_1} \cup \dots \cup X_{i_t} \cup X_{n+1}$), and
- for each $\{i_1, \dots, i_t\} \notin \mathcal{I}'$ there is a dependent set that contains a singleton precisely when it is contained in one of $X_{i_1}, \dots, X_{i_t}, X_{n+1}$ (again, it is the set $X_{i_1} \cup \dots \cup X_{i_t} \cup X_{n+1}$).

This can be explicitly written in MS_0 , see [27], to give us:

Lemma 3.1.3. *Let N be a matroid. There is a sentence ϕ_N in MS_0 such that a matroid M has minor isomorphic to N precisely when M satisfies ϕ_N .*

This is useful as it means that MS_0 is sufficient to define classes with finitely many obstructions. In particular, we can MS_0 -axiomatize \mathbb{F} -representability when \mathbb{F} is a finite field [13].

3.1.3 Ingleton's inequality

Note that minors provide an infinite axiomatization of representability: we simply have an axiom for every excluded minor. However it is conceivable that these could be replaced by a finite set of necessary and sufficient conditions for real-representability. Indeed, there are necessary conditions for representability that rule out infinitely many excluded minors for real-representability.

For example, Ingleton observed an inclusion-exclusion principle on the dimensions of the intersection of two subspaces with relation to another two subspaces [17]. For four subspaces A, B, C, D of a vector space, the dimension of the intersection of A and B is less than or equal to the sum of the dimension-deficit of C in their intersection, the dimension-deficit of D in their intersection, and the dimension of the intersection of C and D :

$$\dim(A \cap B) \leq [\dim(A \cap B) - \dim(A \cap B \cap C)] + [\dim(A \cap B) - \dim(A \cap B \cap D)] + \dim(C \cap D)$$

We can rephrase this as follows in terms of bases A_0, B_0, C_0, D_0 of A, B, C, D and to allow a conversion to MS_0 :

For all independent subsets A_0, B_0, C_0, D_0 and each independent subset S for which A_0 and B_0 are maximal independent sets in $A_0 \cup S$ and $B_0 \cup S$, respectively, there exist independent sets X, Y, Z whose union contains S and such that:

- $X \cup C_0$ is independent and A_0 and B_0 are maximal independent sets in $A_0 \cup X$ and $B_0 \cup X$, respectively,
- $Y \cup D_0$ is independent and A_0 and B_0 are maximal independent sets in $A_0 \cup Y$ and $B_0 \cup Y$, respectively,
and
- C_0 and D_0 are maximally independent sets in $C_0 \cup Z$ and $D_0 \cup Z$, respectively.

Ingleton's inequality rules out infinitely many obstructions to real representability [28]. One might hope that we only need finitely many more conditions to axiomatize representability. How can we show that this is not the case?

3.2 Ineffable properties

We now need a way to show that a property P is not finitely axiomatizable. The strategy presented is similar to Myhill and Nerode's characterization of regular languages [33] and was used by Mayhew, Newman and Whittle. We find two matroids that should satisfy the same MS_0 -sentences on a given number of variables, yet have one that has the property P and one that does not. First, however, we need a better understanding of what we can say about matroids using finitely many variables.

3.2.1 Depth- k truth tables

We introduce *depth- k truth tables* as a data structure that captures what we can say about a pair (E, \mathcal{I}) using k quantifiers. For a positive integer k , we define depth- k truth tables recursively.

For $A_1, \dots, A_n \subseteq E$, the *depth-0 truth table* $\mathcal{T}_0(E, \mathcal{I}; A_1, \dots, A_n)$ of a pair (E, \mathcal{I}) is a $(n + 2) \times n$ -table whose rows are indexed by $X_1, \dots, X_n, \text{Ind}, \text{Sing}$ and whose columns are indexed by X_1, \dots, X_n . It is essentially a truth table for *atomic expressions*, that is, expressions with no connectives or quantifiers. The (X_i, X_j) -entry is \top if $A_i \subseteq A_j$ and \perp otherwise and the (Ind, X_j) -entry is \top if A_j is independent and \perp otherwise. However, the (Sing, X_j) -entry is defined slightly differently: it is \top if A_j is a singleton, but it takes on a value of 0 if A_j is empty and \perp if A_j has at least two elements. Note that we get the same depth-0 truth table for different (E, \mathcal{I}) and A_1, \dots, A_n precisely when the the same atomic expressions are satisfied and the empty sets remain the same.

The following illustrates the key property of depth-0 truth tables. Suppose we are given a quantifier-free formula $\phi(X_1, \dots, X_k)$. We may use the depth-0 truth table $\mathcal{T}_0(E, \mathcal{I}; A_1, \dots, A_n)$ to determine whether or not the sets A_1, \dots, A_n satisfy $\phi(X_1, \dots, X_n)$ according to (E, \mathcal{I}) . Indeed, the table gives us the truth value of all atomic expressions that may occur as constituents and we may use boolean arithmetic to calculate the result of evaluating $\phi^{(E, \mathcal{I})}(A_1, \dots, A_n)$. We now adapt truth tables to allow us to consider formulae with a fixed number of quantified variables.

For a positive integer k , and $A_1, \dots, A_n \subseteq E$, we define the *depth- k truth table* $\mathcal{T}_k(E, \mathcal{I}; A_1, \dots, A_n)$ of a pair (E, \mathcal{I}) as the set of all the depth- $(k - 1)$ truth tables that are possible in (E, \mathcal{I}) with the given A_1, \dots, A_n and an additional parameter $A_{n+1} \subseteq E$; that is

$$\mathcal{T}_k(E, \mathcal{I}; A_1, \dots, A_n) := \{\mathcal{T}_{k-1}(E, \mathcal{I}; A_1, \dots, A_n, A_{n+1}) : A_{n+1} \subseteq E\}.$$

Note that truth tables with depth at least 1 are the same precisely when they have the same elements, regardless of the choice of (E, \mathcal{I}) and A_1, \dots, A_n .

Suppose that we are given a formula $\phi(X_1, \dots, X_n)$ with an additional k variables X_{n+1}, \dots, X_{n+k} that are quantified. By standard conversion rules, we may assume that ϕ is of the *prenex normal form* — with all quantifiers and their variables occurring as a prefix. This implies that ϕ is of the form $\forall X_{n+1}\psi$ or $\exists X_{n+1}\psi$, for a formula $\psi(X_1, \dots, X_n, X_{n+1})$ with quantifier depth $k - 1$. As an induction hypothesis, we assume that the depth- $(k - 1)$ truth table $\mathcal{T}_{k-1}(E, \mathcal{I}; A_1, \dots, A_n, A_{n+1})$ is sufficient to determine $\psi^{(E, \mathcal{I})}(A_1, \dots, A_n, A_{n+1})$, the truth value of ψ interpreted for (E, \mathcal{I}) with X_i assigned to A_i for i from 1 to $n + 1$. However, the depth- k truth table $\mathcal{T}_k(E, \mathcal{I}; A_1, \dots, A_n)$ contains all the possible truth tables arising from different choices of A_{n+1} . Because of this, $\mathcal{T}_k(E, \mathcal{I}; A_1, \dots, A_n)$ is then sufficient to determine the truth value of ϕ for X_i assigned to A_i for i from 1 to n . Specifically, if ϕ is of the form $\forall X_{n+1}\psi$, then $\mathcal{T}_k(E, \mathcal{I}; A_1, \dots, A_n)$ determines a value of “true” precisely when all of its elements determine a value of “true” for ψ . Similarly, if ϕ is of the form $\exists X_{n+1}\psi$, then $\mathcal{T}_k(E, \mathcal{I}; A_1, \dots, A_n)$ determines a value of “true” precisely when one of its elements determine a value of “true” for ψ .

We have also just given the idea of the proof of the following theorem.

Theorem 3.2.1 (Mayhew, Newman, Whittle [27]). *To determine whether or not (E, \mathcal{I}) satisfies a k -variable sentence, it is enough to have the zero-parameter depth- k truth table $\mathcal{T}_k(E, \mathcal{I})$.*

It is not difficult to check the following by inducting on k .

Remark 3.2.2 (Mayhew, Newman, Whittle [27]). *There are finitely many depth- k truth tables on zero parameters.*

It is also useful to note that the zero-parameter depth- k truth table of two matroids determines the zero-parameter depth- k truth table of their direct sum:

Lemma 3.2.3 (Mayhew, Newman, Whittle [27]). *Given zero-parameter depth- k truth tables \mathcal{T} and \mathcal{T}' , there is a unique zero-parameter depth- k truth table $\mathcal{T} \oplus \mathcal{T}'$ such that if M and M' have zero-parameter depth- k truth table \mathcal{T} and \mathcal{T}' respectively, then $M \oplus M'$ has zero-parameter depth- k truth table $\mathcal{T} \oplus \mathcal{T}'$.*

3.2.2 k -equivalence

We now introduce equivalence relations that give a notion of which matroids are indistinguishable when we limit our descriptive power to a given number of variables. For a

positive integer k , we say that M_1 and M_2 are k -equivalent when M_1 and M_2 have the same zero-parameter depth- k truth table.

Note that as there are finitely many depth- k truth tables on zero parameters (Lemma 3.2.2),

Remark 3.2.4 (Mayhew, Newman, Whittle [27, Lemma 1.3]). *There are finitely many k -equivalence classes for a given k .*

As the depth- k truth table of a direct sum is determined by the depth- k truth tables of its summands (see Lemma 3.2.3), we could have defined k -equivalence as follows:

Remark 3.2.5. *M_1 and M_2 are k -equivalent when, for any sentence ϕ in MS_0 of quantifier depth at most k and any matroid N with $E(N) \cap E(M_1) = E(N) \cap E(M_2) = \emptyset$, we have that $M_1 \oplus N$ satisfies ϕ if and only if $M_2 \oplus N$ satisfies ϕ .*

This phrasing is more useful as we will have a sentence ϕ and matroid N as a certificate of inequivalence. When stated as above, it is also clear that k -equivalence is analogous to the equivalence on strings used by Myhill and Nerode [33].

3.3 The non-axiomatizability of representability

We now present Mayhew, Newman, and Whittle's lovely proof that representability is not axiomatizable in MS_0 [27].

Theorem 3.3.1 (Mayhew, Newman, Whittle [27, Theorem 1.1]). *There is no sentence ϕ in MS_0 such that a matroid is representable precisely when it satisfies ϕ .*

Proof. Suppose that a sentence ϕ characterizes representability. Say ϕ has k variables. There are infinitely many primes, but only finitely many equivalence classes for k -equivalence by Lemma 3.2.4. Thus we have two projective planes $M = \text{PG}(2, p)$ and $M' = \text{PG}(2, p')$ in the same equivalence class for distinct primes p and p' . However, $M \oplus M$ is a representable matroid, while $M' \oplus M$ is not. So ϕ and $N = M = \text{PG}(2, p)$ certify that M and M' are not in the same equivalence class, our contradiction. \square

3.4 The non-axiomatizability of real-representability

The proof that representability is not finitely axiomatizable (Theorem 3.3.1) required that representability not be closed under direct sums. However, representability over a fixed field is closed under direct sums; we need another operation to take its place. This will be proper amalgamation (see Section 2.3). With appropriate alterations to the previous technique, Mayhew, Newman, and Whittle proved the following:

Theorem 3.4.1 (Mayhew, Newman, Whittle [27, Theorem 1.2]). *Let \mathbb{F} be an infinite field. There is no sentence $\phi_{\mathbb{F}}$ in MS_0 such that a matroid is \mathbb{F} -representable precisely when it satisfies $\phi_{\mathbb{F}}$.*

We will use the same tools they introduce but different constructions to prove the following, slightly stronger result for the case where $\mathbb{F} = \mathbb{R}$:

Theorem 3.4.2 (Campbell). *There is no sentence ϕ in MS_0 such that a complex-representable matroid is real-representable precisely when it satisfies ϕ .*

3.4.1 Alterations to the previous technique

We now look at how we must modify k -equivalence and depth- k truth tables to deal with proper amalgams across a fixed matroid. Here we restrict to amalgams across a fixed line, but this can be done more generally.

We first adapt the equivalence relation; to do so, we draw inspiration from Remark 3.2.5 (and Myhill and Nerode [33]). Fix a matroid R consisting of a single line. Let \mathcal{M}_R be the set of all matroids that contain R as a restriction. For a positive integer k , we say that M_1 and M_2 in \mathcal{M}_R are $(k; R)$ -equivalent when, for any k -variable sentence ϕ and any matroid N in \mathcal{M}_R with $E(N) \cap E(M_1) = E(N) \cap E(M_2) = E(R)$, we have that the proper amalgam $M_1 \oplus_R N$ satisfies ϕ if and only if $M_2 \oplus_R N$ satisfies ϕ .

We now modify depth- k truth tables so that for a given M in \mathcal{M}_R there is enough information to determine whether an amalgam $M \oplus_R N$ satisfies a k -variable sentence ϕ . While quantifiers behave as before, we must include additional information in the depth-0 truth tables. To be able to determine the independent sets in the proper amalgam $M \oplus_R N$, we need to know how the subsets of $E(M)$ interact with R . Specifically, for a subset A_j of $E(M)$, we need to know whether A_j is dependent, but otherwise the elements in $A_j \cap E(R)$ and the elements of R in the span of $A_j - E(R)$. We can do this by simply including this information: the (Ind, X_j) -entry of the depth-0 truth table $\mathcal{T}_0(M; A_1, \dots, A_n)$ is \perp

if A_j is dependent and the ordered pair $(A_j \cap E(R), \text{cl}(A_j - E(R)) \cap E(R)) \in E(R) \times E(R)$ otherwise. With this additional information included in each entry on each Ind-row, Mayhew, Newman, and Whittle proved the following analogue of Lemma 3.2.4.

Lemma 3.4.3 (Mayhew, Newman, Whittle [27, Lemma 1.4]). *For each integer k , there are only finitely many equivalence classes for $(k; R)$ -equivalence.*

3.4.2 New constructions

We now consider the constructions that will take the place of the projective geometries from the previous section. Here we differ from the proof of Mayhew, Newman, and Whittle [27].

For an odd prime p , let M_p be the rank- $(p + 1)$ swirl-like matroid given by the $\mathbb{Q}(x)$ -matrix

$$A_p(x) = \begin{pmatrix} b_0 & [1]_0 & [x^{-1}]_0 & \dots & [x^{-p}]_0 & b_1 & [1]_1 & [x]_1 & [1]_2 & [x]_2 & \dots & [1]_p & [x]_p \\ \hline 1 & 1 & 1 & \dots & 1 & 0 & 0 & 0 & 0 & 0 & \dots & -1 & -x \\ 0 & -1 & -x^{-1} & \dots & -x^{-p} & 1 & 1 & 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & -1 & -x & 1 & 1 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & -1 & -x & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 \end{pmatrix}$$

Note that the labellings are encoded with respect to the transversal $\{[1]_0, \dots, [1]_p\}$.

The dependent transversals of M_p are those whose labels multiply to 1, as in Section 2.2.2. This enforces representations that encode this labelling.

Let $z = x^p$, so we can also label the element $[x^{-p}]_0$ as $[z^{-1}]_0$. Now let R be the line restriction on the set $\{b_0, [1]_0, [z^{-1}]_0, b_1\}$. So that we may consider an amalgam across R , let M'_p be a copy of M_p where all the elements not in R have been renamed and where the indeterminant x has been replaced with y in labellings.

Intuitively, when we consider the matroid $M_q \oplus_R M'_p$ for odd primes p and q , the encodings enforce $x^q = z = y^p$ with $x \neq y$. This leads to the following two propositions.

Proposition 3.4.4. *If q and p are distinct odd primes, then $M_q \oplus_R M'_p$ is real-representable.*

Proof. Let γ be a positive transcendental real number. Let $\alpha = \gamma^{\frac{1}{q}}$ and $\beta = \gamma^{\frac{1}{p}}$. Consider the homomorphism from $\mathbb{Q}(z, x, y)$ to \mathbb{R} by evaluating z, x, y as γ, α, β , respectively.

Let L_q and L'_p be the closures of $E(R)$ in M_q and M'_p respectively. As q and p are distinct primes, L_q and L'_p have no coincident values labelled besides those in $E(R)$. Thus by Lemma 2.3.4, the proper amalgam $M_q \oplus_R M'_p$ is real-representable. \square

Proposition 3.4.5. *For any odd prime p , the amalgam $M_p \oplus_L M'_p$ is complex-representable but not real-representable.*

Proof. Let γ be a transcendental complex number, and ω be a primitive p^{th} root of unity. Let $\alpha = \gamma^{\frac{1}{p}}$ and $\beta = \gamma\omega$. Consider the homomorphism from $\mathbb{Q}(z, x, y)$ to \mathbb{C} by evaluating z, x, y as γ, α, β , respectively.

Let L_p and L'_p be the closures of $E(R)$ in M_p and M'_p respectively. As ω is a primitive root of unity, L_p and L'_p have no coincident values besides those in $E(R)$. Thus by Lemma 2.3.4, the proper amalgam $M_p \oplus_R M'_p$ is complex-representable.

Suppose that $M_p \oplus_R M'_p$ has a real-representation. Say z is evaluated as γ on the line $E(R)$. The representation of $M_p \oplus_R M'_p$ induces encodings on the swirl-like matroid restrictions M_p and M'_p . Say x, y are evaluated as α, β in M_p and M'_p respectively. The encoding of $E(R)$ in the restriction M_p enforces that α is the p^{th} root of γ . However, the restriction M'_p enforces that the p^{th} root of γ is β . However, as p is odd, γ only has one real p^{th} root by distribution of the absolute value over products. Thus $M_p \oplus_R M'_p$ is not real-representable. \square

3.4.3 Analogous proof

Proof of Theorem 3.4.2: Suppose that we have a sentence ϕ in MS_0 that characterizes real-representability for complex-representable matroids. Say ϕ has k variables. There are infinitely many primes, but only finitely many equivalence classes for $(k; R)$ -equivalence by Lemma 3.4.3. Thus we have two swirls M_p and M_q in the same equivalence class for distinct primes p and q . Note that both $M_q \oplus_R M'_p$ and $M_p \oplus_R M'_q$ are complex-representable by Propositions 3.4.4 and 3.4.5. However, by these same propositions, $M_q \oplus_R M'_p$ is real-representable, while $M_p \oplus_R M'_p$ is not. So ϕ and $N = M'_p$ certify that M_p and M_q are not in the same equivalence class for $(k; R)$ -equivalence, a contradiction. This means, that we cannot have such a sentence ϕ in MS_0 that characterizes real-representability for complex-representable matroids. \square

Chapter 4

Excluded minors

This section is based on [8] co-authored with Jim Geelen and published in the Journal of Combinatorial Theory, Series B.

Instead of inspecting why a matroid is real-representable, it may be easier to say why a matroid is **not** real-representable. This has turned out to be the case for representability over a finite field. Indeed, Tutte showed that a matroid is not $\text{GF}(2)$ -representable if and only if it contains the 4-point line as a minor [41]. Geelen, Gerards, and Whittle announced a proof of Rota's Conjecture [13], that is, for any finite field \mathbb{F} , there is a finite collection of matroids, $\mathcal{E}_{\mathbb{F}}$, such that a matroid is non- \mathbb{F} -representable precisely when it contains a member of $\mathcal{E}_{\mathbb{F}}$ as a minor.

4.0.1 Excluded minors for real-representability

We consider characterizing the set of excluded minors for the class of real-representable matroids. In contrast to the case for finite fields, Lazarsen [22, Theorem 1] showed that there are infinitely many excluded minors for real-representability. This in itself does not preclude the possibility of a simple structural description. For example, Bonin [3, Theorem 3.1] described the excluded minors for lattice-path matroids, despite the fact that the list is infinite.

Mayhew, Newman, and Whittle [28] have effectively settled the matter by proving the following striking result.

Theorem 4.0.1 (Mayhew, Newman, Whittle [28]). *For any infinite field \mathbb{F} , each \mathbb{F} -representable matroid is a minor of an excluded minor for \mathbb{F} -representability.*

This essentially implies that the excluded minors are at least as structurally complicated as the real-representable matroids themselves. This is less surprising when we consider that non-representable matroids are rather wild in comparison to representable matroids. For instance, Nelson showed that asymptotically almost all matroids are non-representable [32]. Thus it is more unexpected that the same issue arises even if we only try to describe the representable excluded minors for real-representability: we see that even the complex-representable excluded minors are at least as wild as the class of real-representable matroids.

Theorem 4.0.2 (Campbell, Geelen [8]). *Each real-representable matroid is a minor of a complex-representable excluded minor for real-representability.*

Given a real-representable matroid, M , we wish to construct an excluded minor for real-representability, \mathcal{X}_M , that is complex-representable and contains M as a minor. We will do this by combining a known complex-representable excluded minor, N , with M using complexity-preserving operations.

4.0.2 Complex-representability preserving operations

For a field \mathbb{F} , it is well-known and easy to show that the class of \mathbb{F} -representable matroids is closed under isomorphisms, minors, adding coloops, and direct sums. Direct sums (and adding coloops) allow us to make a “foundation” and build “scaffolding” for our construction, while isomorphism and minors allow us to manipulate and get rid of auxiliary structure we use during construction. However, we need a method to “fill in” and add structure to our construction. We will use principal extensions for this.

Principal Extension

Let F be a flat of a matroid M . A *principal extension of M into the flat F* is the matroid M' on a ground set $E(M) \cup \{e\}$ where $M' \setminus e = M$ and a subset of $E(M)$ spans e if and only if it spans F . We say that M' is obtained by *freely placing e in F* , and *freely placing e* when $F = E(M)$.

It is standard material that principal extension is well defined [34, Proposition 7.2.6] and preserves \mathbb{F} -representability for all infinite fields \mathbb{F} , see [28, Lemma 2.1].

M -constructed matroids

For a matroid M , we say that a matroid M' is M -constructed if it can be obtained from M by a sequence of the following operations: renaming elements, deletion, contraction, adding coloops, and principal extensions. Consequently, if M is \mathbb{F} -representable for an infinite field \mathbb{F} and M' is an M -constructed matroid, then M' is also \mathbb{F} -representable

Let e be an element of a matroid M . The *series extension* of e in M is the matroid M' obtained by coextending M by an element e' so that $\{e, e'\}$ is a series pair.

Lemma 4.0.3. *Let e be an element of a matroid M . If M' is the series extension of e in M , then M' is M -constructed.*

Proof. Let M_1 be obtained from M by adding a coloop e' and then freely placing an element e'' in the flat spanned by $\{e, e'\}$. Then M' is obtained from $M_1 \setminus e$ by renaming e'' as e . \square

4.1 Natural classes

As the construction we will consider only requires the aforementioned operations, we instead prove generalization of Theorem 4.0.2. Similarly, with inspired by Mayhew, Newman, and Whittle's construction, Matúš explored classes for which a generalization of Theorem 4.0.1 holds [26].

We say that a class \mathcal{M} is *natural*, when it is non-empty and closed under isomorphism, minors, adding coloops, direct sums, and principal extensions. Consequently, if M is a matroid in a natural class \mathcal{M} and M' is an M -constructed matroid, then M' is contained in \mathcal{M} . So by Lemma 4.0.3, natural classes are closed under series extensions.

Recall that a *pinned extension into a 3-separation* includes a new element that is “pinned” on either side of the 3-separation, see Section 2.3.1. We prove the following generalization of Theorem 4.0.2.

Theorem 4.1.1 (Campbell, Geelen [8]). *Let \mathcal{M} and \mathcal{N} be natural classes where $\mathcal{M} \subsetneq \mathcal{N}$. If \mathcal{M} is closed under pinned extensions into 3-separations, then each matroid in \mathcal{M} is a minor of an excluded minor of \mathcal{M} that is also in \mathcal{N} .*

Note that the class of all matroids is natural. Furthermore, it follows directly from the definition that arbitrary intersections of natural classes are also natural. Thus, under the

subset relation, natural classes of matroids form a lattice. This is also true for classes that are closed under pinned extensions into 3-separations.

We already know that for an infinite field \mathbb{F} , the \mathbb{F} -representable matroids form a natural class that is closed under pinned extensions into 3-separations, see Sections 4.0.2 and Lemma 2.3.6, respectively. The “algebraic matroids” for a fixed field can also be shown to be a natural class, see [34, Corollary 6.7.14;26, Lemma 13]. As we will use in Section 6, the class of “orientable matroids” is natural [1, page 330]. Of particular interest to us is the class of “gammoids”, as in Corollary 4.1.3 we will see that this is the minimal natural class.

There are certainly interesting classes that are not “natural”. For a prime power q , the $\text{GF}(q)$ -representable matroids are not a natural class: the uniform matroid $U_{2,q+1}$ is $\text{GF}(q)$ -representable while the principal extension $U_{2,q+2}$ is not. More generally, as the class of gammoids is the minimal natural class, any class of matroids that does not contain all gammoids is not natural, regardless of how basic it is.

4.1.1 Gammoids

Let H be a bipartite graph whose vertices have bipartition $A \sqcup B$. Let \mathcal{I} be the set of subsets of A that can be covered by a matching in H . Edmonds and Fulkerson showed that $T = (A, \mathcal{I})$ is a matroid [9], and named such matroids *transversal matroids*. Brylawski showed that this transversal matroid $T = (A, \mathcal{I})$ can be real-represented by labelling the vertices of the standard basis vector simplex in $\mathbb{R}^{|B|}$ with B and mapping each $a \in A$ to a freely placed point on the face spanned by its neighbours in H [5, Theorem 3.1 and Corollary 3.1] (see also Oxley [34, Proposition 11.2.26]). Recall that while B was again used for this description, it is not part of the ground set of the matroid $T = (A, \mathcal{I})$.

Here we define a *gammoid* as a matroid that can be obtained from a transversal matroid with a sequence of contractions. Gammoids are usually defined by the connectivity function to a fixed set in a digraph, as in Oxley [34, page 97 and 109], and the equivalence of these definitions was shown by Ingleton and Piff [18, Theorem 3.5].

We now show that the class of gammoids is the minimal class that is closed under isomorphisms, minors, adding coloops, and principal extensions. We can restate this as the follows.

Theorem 4.1.2. *A matroid is a gammoid if and only if it is $U_{0,0}$ -constructible.*

Proof. As an immediate consequence of Brylawski's description above, all transversal matroids are $U_{0,0}$ -constructible. Next, by Ingleton and Piff description of the class of gammoids, all gammoids are $U_{0,0}$ -constructible.

Another consequence of the aforementioned descriptions is that gammoids are closed under isomorphisms, minors, and adding coloops. It remains to be shown that the class of gammoids is closed under principal extensions. For a gammoid G with ground set E , say G is the contraction of a transversal matroid T by a set $X \subseteq E(T)$. Consider T represented by points S in \mathbb{R}^r which are freely placed on the faces of the simplex whose vertex set $B \subseteq \mathbb{R}^r$ is the standard basis. Suppose we have $e \in E$ that does not lie on a vertex of this simplex, but lies in the affine span of $B_e \subseteq B$. We coextend T , by turning e into a series pair $\{x, e\}$ to get T' as in Lemma 4.0.3. By embedding \mathbb{R}^r in \mathbb{R}^{r+1} , we have that T' is represented by points that lie in the faces of the standard simplex with vertex set $B \cup \{e\}$; the points in $S - \{e\}$ lie in the span of B as before, but x lies in the face spanned by $B_e \cup \{e\}$. Note $T'/X \cup \{x\} = T/X = G$ but now e is on a vertex of the simplex in the representation. In this way we may assume that all of E lies on the standard basis of a representation of T . Consider a principal extension G_F into a flat $F \subseteq E$. Note that the principal extension T_F into $F \subseteq E$ is also a transversal matroid as F are the vertices of a face. Thus $T_F/X = G_F$ is a gammoid, as we wanted to show. \square

As it is also easy to show that the class of gammoids is closed under direct sums, we have the following.

Corollary 4.1.3. *The class of gammoids is the minimal natural class.*

In a matroid N , we say an element p is *freer* than an element q , when every subset of $E(N) - \{p, q\}$ that spans p also spans q . A pair of elements $\{p, q\}$ is *incomparable* when there exist subsets Y_p and Y_q of $E(N) - \{p, q\}$ such that Y_p spans p but not q while Y_q spans q but not p .

Lemma 4.1.4. *Matroids with no incomparable pair are gammoids.*

Proof. Suppose that N has no incomparable pair. Then there is an ordering (e_1, \dots, e_n) of $E(N)$ such that e_j is freer than e_i whenever $1 \leq i < j \leq n$. So either N is the empty matroid and hence $U_{0,0}$ -constructible, or else N has a freest element e_n . Now either e_n is a coloop in N or N is obtained by placing e_n freely in $N \setminus e_n$. Note that $N \setminus e_n$ has no incomparable pair so we may inductively assume that $N \setminus e_n$ is $U_{0,0}$ -constructible, and, hence, N is $U_{0,0}$ -constructible. \square

4.2 Constructing a “major” excluded minor

Our construction for the proof of Theorem 4.1.1 is based on a known excluded minor N contained in \mathcal{N} . However, we cannot take an arbitrary excluded minor for the class \mathcal{M} ; we require that N has a “special” pair of elements $\{p, q\}$. We see that there exist such an N in Section 4.2.1. In Section 4.2.2, we see that given a matroid M in \mathcal{M} , we do not lose generality by assuming that the ground set of M can be partitioned into two bases A and B : otherwise M is the minor of such a matroid in \mathcal{M} . Finally, in Section 4.2.3 we see how we can modify $M \oplus N$ so that we have collections of elements “pointing at” p and q in a copy of N , and how this gives us our desired excluded minor.

4.2.1 Picking a base excluded minor N in \mathcal{N}

If N_1 and N_2 are matroids on a common ground set E , then we say that N_2 is *freer* than N_1 if $r_{N_2}(X) \geq r_{N_1}(X)$ for each subset X of E . Let p and q be distinct elements of a matroid N and let N' denote the matroid obtained from N by freely adding a new element e into the flat spanned by $\{p, q\}$. We denote by $N_{p \rightarrow q}$ the matroid obtained from $N \wedge p$ by renaming e as p . Note that e is freer than p in N' and so any subset X spanning e would also span p in N' . Hence $r_{N_{p \rightarrow q}}(X) \geq r_N(X)$ for each subset X of E and so $N_{p \rightarrow q}$ is freer than N .

Lemma 4.2.1. *Let \mathcal{M} and \mathcal{N} be natural classes of matroids. If $\mathcal{M} \subsetneq \mathcal{N}$, then there is an excluded minor N for \mathcal{M} in \mathcal{N} with a pair $\{p, q\}$ of incomparable elements such that $N_{p \rightarrow q}$ and $N_{q \rightarrow p}$ are both contained in \mathcal{M} .*

Proof. Since $\mathcal{M} \subsetneq \mathcal{N}$, there is an excluded minor for \mathcal{M} in \mathcal{N} . Among all excluded minors for \mathcal{M} in \mathcal{N} we choose N satisfying:

- $|E(N)|$ is minimum, and
- subject to this, N is *freest* with ground set $E(N)$ (that is, there is no other excluded minor N' with ground set $E(N)$ that is freer than N).

By Theorem 4.1.2, \mathcal{M} contains all gammoids and, by Lemma 4.1.4, N has an incomparable pair $\{p, q\}$. Note that $N_{p \rightarrow q}$ and $N_{q \rightarrow p}$ are both N -constructed and hence they are both contained in \mathcal{N} . Moreover, both $N_{p \rightarrow q}$ and $N_{q \rightarrow p}$ are freer than N . So, by our choice of N , both $N_{p \rightarrow q}$ and $N_{q \rightarrow p}$ are contained in \mathcal{M} , as required. \square

4.2.2 Preprocessing M ; Bipartition by bases

The following result is essentially due to Mayhew, Newman, and Whittle [28, Lemma 2.2].

Lemma 4.2.2. *For any matroid M , there is an M -constructed matroid M' such that M' has an M -minor and the ground-set of M' can be partitioned into two bases.*

Proof. Let A_0 and B_0 denote two $r(M)$ -element sets that are disjoint from $E(M)$ and from each other. We extend M by adding the elements $A_0 \cup B_0$ freely to obtain the matroid M_1 . Note that the only circuits containing members of A_0 or B_0 are spanning circuits and so A_0 and B_0 are bases. Next, we construct M' from M_1 by a sequence of series extensions for each element in $E(M)$; we rename the elements so that, for each $e \in E(M)$, the corresponding series-pair in M' is $\{e_1, e_2\}$. Note M' has bases $A_1 = A_0 \cup \{e_1 : e \in E(M)\}$ and $B_1 = B_0 \cup \{e_2 : e \in E(M)\}$ which partition $E(M')$, as required. \square

Recall that we are trying to find an excluded minor for \mathcal{M} that contains M as a minor. By Lemma 4.2.2, we lose no generality in assuming that M can be partitioned into two bases. To prove Theorem 4.0.1, Mayhew, Newman, and Whittle use a similar assumption to construct a matroid with a circuit-hyperplane whose relaxation does not lose M as a minor but creates a violation of Ingleton's inequality (similar to the Vámos matroid). As we would like our excluded minor to be in \mathcal{N} , we will instead use the bases given by Lemma 4.2.2 to “point at” the elements p and q in the excluded minor N given by Lemma 4.2.1.

4.2.3 Using M to cannibalize N

In the construction and in all subsequent results in this section,

- N and M are matroids with disjoint ground sets,
- (A, B) is a partition of $E(M)$ into two bases, and
- p and q are distinct elements of N .

We build an $(M \oplus N)$ -constructed matroid $\mathcal{X}(N, p, q; M, A, B)$ as follows. The input and output of the construction process are depicted in Figure 4.1.

We first build a matroid $\mathcal{X}_1(N, p, q; M, A, B)$ from $M \oplus N$ by freely placing elements a and b in the flats spanned by $E(M) \cup \{p\}$ and $E(M) \cup \{q\}$ respectively; then, for each

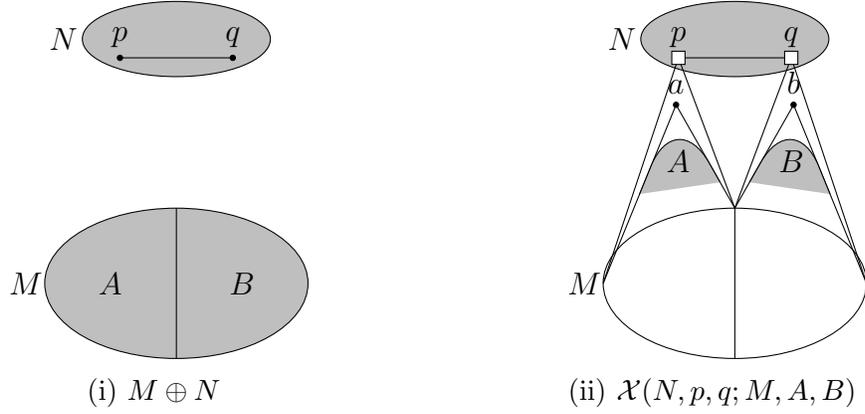


Figure 4.1: The input and output of the construction process.

$x \in A$, we freely place an element x_a in $\{x, a\}$, and, for each $y \in B$, we freely place an element y_b in $\{y, b\}$.

We then obtain $\mathcal{X}_2(N, p, q; M, A, B)$ from $\mathcal{X}_1(N, p, q; M, A, B)$ as follows: for each $x \in A$ and $y \in B$, we delete x and y and rename x_a and y_b as x and y respectively. Finally, we let $\mathcal{X}(N, p, q; M, A, B) = \mathcal{X}_2(N, p, q; M, A, B) \setminus \{p, q\}$. Note that if $\{p, q\}$ is an independent pair in N , then $\mathcal{X}(N, p, q; M, A, B)$ contains M as the minor $\mathcal{X}(N, p, q; M, A, B) / \{a, b\} \setminus E(N)$.

The main result in this section is the following.

Theorem 4.2.3. *Let \mathcal{M} be a natural class that is closed under pinned extensions into 3-separations. If*

- (i) N is an excluded minor for \mathcal{M} ,
- (ii) p and q are an incomparable pair of elements in N such that $N_{p \rightarrow q}$ and $N_{q \rightarrow p}$ are both in \mathcal{M} ,
- (iii) M is in \mathcal{M} , and
- (iv) (A, B) is a partition of $E(M)$ into bases,

then $\mathcal{X}(N, p, q; M, A, B)$ is an excluded minor for the class \mathcal{M} .

$\mathcal{X}(N, p, q; M, A, B)$ is not in \mathcal{M}

We will start by proving that $\mathcal{X}(N, p, q; M, A, B)$ is not in \mathcal{M} . For this we require the following results. The first of these results gives us some simple structural properties of $\mathcal{X}(N, p, q; M, A, B)$.

Lemma 4.2.4. *Let $\mathcal{X} = \mathcal{X}(N, p, q; M, A, B)$. If $\{p, q\}$ is independent and coindependent in N , then $(E(N) - \{p, q\}, E(M) \cup \{a, b\})$ is a 3-separation in \mathcal{X} and $\square_{\mathcal{X}}(A \cup \{a\}, E(N) - \{p, q\}) = \square_{\mathcal{X}}(B \cup \{b\}, E(N) - \{p, q\}) = 1$.*

Proof. Let $\mathcal{X} = \mathcal{X}(N, p, q; M, A, B)$ and $\mathcal{X}_2 = \mathcal{X}_2(N, p, q; M, A, B)$. Note that $\square_{\mathcal{X}_2}(E(N), E(M) \cup \{a, b\}) = 2$ and $\square_{\mathcal{X}_2}(A \cup \{a\}, E(N)) = \square_{\mathcal{X}_2}(B \cup \{b\}, E(N)) = 1$. Then, since $\{p, q\}$ is coindependent in N , we have that $(E(N) - \{p, q\}, E(M) \cup \{a, b\})$ is a 3-separation in \mathcal{X} and $\square_{\mathcal{X}}(A \cup \{a\}, E(N) - \{p, q\}) = \square_{\mathcal{X}}(B \cup \{b\}, E(N) - \{p, q\}) = 1$. \square

The following result shows that, if we extend $\mathcal{X}(N, p, q; M, A, B)$ by nonloop elements p and q such that p is spanned by both $A \cup \{a\}$ and $E(N) - \{p, q\}$ whereas q is spanned by both $B \cup \{b\}$ and $E(N) - \{p, q\}$, then we retrieve $\mathcal{X}_2(N, p, q; M, A, B)$.

Lemma 4.2.5. *Let \mathcal{X}' be an extension of $\mathcal{X}(N, p, q; M, A, B)$ by nonloop elements p and q such that p is spanned by both $A \cup \{a\}$ and $E(N) - \{p, q\}$ whereas q is spanned by both $B \cup \{b\}$ and $E(N) - \{p, q\}$. If $\{p, q\}$ is an incomparable pair in N , then $\mathcal{X}' = \mathcal{X}_2(N, p, q; M, A, B)$.*

Proof. Let $\mathcal{X} = \mathcal{X}(N, p, q; M, A, B)$ and $\mathcal{X}_2 = \mathcal{X}_2(N, p, q; M, A, B)$. Since $\{p, q\}$ is an incomparable pair in N , $\{p, q\}$ is both independent and coindependent and there exist sets $Y_p, Y_q \subseteq E(N) - \{p, q\}$ such that Y_p spans p but not q and Y_q spans q but not p . Note that $\square_{\mathcal{X}}(Y_p, E(M) \cup \{a, b\}) = \square_{\mathcal{X}}(Y_p, A \cup \{a\}) = 1$, and $\square_{\mathcal{X}}(Y_q, E(M) \cup \{a, b\}) = \square_{\mathcal{X}}(Y_q, A \cup \{a\}) = 1$. Combining these with Lemma 4.2.4 gives us that $(Y_p, A \cup \{a\})$ pins p and $(Y_q, B \cup \{b\})$ pins q . Moreover $\mathcal{X} \setminus \{p, q\} = \mathcal{X}_2 \setminus \{p, q\} = \mathcal{X}$. As pinned extensions are unique by Lemma 2.3.5, it follows that $\mathcal{X}' = \mathcal{X}_2$. \square

We can now show that $\mathcal{X}(N, p, q; M, A, B)$ is not in \mathcal{M} .

Lemma 4.2.6. *If N is not in the natural class \mathcal{M} and $\{p, q\}$ is an incomparable pair in N , then $\mathcal{X}(N, p, q; M, A, B)$ is not in \mathcal{M} either.*

Proof. Let $\mathcal{X} = \mathcal{X}(N, p, q; M, A, B)$ and $\mathcal{X}_2 = \mathcal{X}_2(N, p, q; M, A, B)$. Since $\{p, q\}$ is an incomparable pair in N , $\{p, q\}$ is both independent and coindependent. We may assume, towards a contradiction, that \mathcal{X} is in \mathcal{M} . By Lemma 4.2.4, there is a pinned extension \mathcal{X}' of

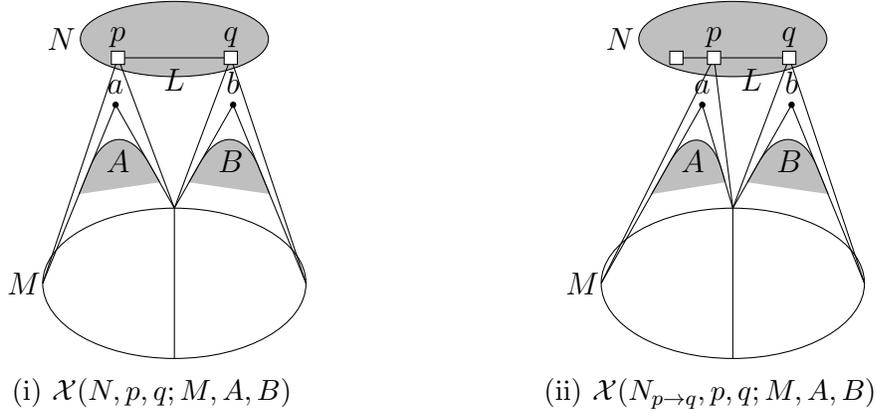


Figure 4.2: Matroids that have the same minor when an element of $A \cup \{a\}$ is deleted or an element of $B \cup \{b\}$ is contracted.

\mathcal{X} by nonloop elements p and q such that p is spanned by both $A \cup \{a\}$ and $E(N) - \{p, q\}$ whereas q is spanned by both $B \cup \{b\}$ and $E(N) - \{p, q\}$. By Lemma 4.2.5, we have $\mathcal{X}' = \mathcal{X}_2$. However, $\mathcal{X}'|E(N) = \mathcal{X}_2|E(N) = N$, which is not in \mathcal{M} . As \mathcal{M} is closed under pinned extensions and deletion, this is a contradiction, which completes the proof. \square

Proper minors of $\mathcal{X}(N, p, q; M, A, B)$ are in \mathcal{M}

It remains to prove that proper minors of $\mathcal{X}(N, p, q; M, A, B)$ are in \mathcal{M} . We do this by showing that every proper minor of $\mathcal{X}(N, p, q; M, A, B)$ is also a minor of one of $\mathcal{X}(N_{p \rightarrow q}, p, q; M, A, B)$, $\mathcal{X}(N_{q \rightarrow p}, p, q; M, A, B)$, or $\mathcal{X}(N', p, q; M, A, B)$, where N' is a proper minor of N . For this we need two preliminary lemmas; the first shows that there is a unique set in $A \cup B \cup \{a, b\}$ that spans p but not q .

Lemma 4.2.7. *If $X \subseteq A \cup B \cup \{a, b\}$ spans p but not q in $\mathcal{X}_2(N, p, q; M, A, B)$, then $X = A \cup \{a\}$.*

Proof. Let M_2 denote the restriction of $\mathcal{X}_2(N, p, q; M, A, B)$ to $A \cup B \cup \{a, b, p, q\}$. We consider an alternate construction of M_2 . Let M_1 be obtained from M by adding coloops p and q and adding a freely to the flat spanned by $A \cup \{p\}$ and b freely to the flat spanned by $B \cup \{q\}$; then, for each $x \in A$, we freely place x_a in $\{x, a\}$ and, for each element $y \in B$, freely place y_b in $\{y, b\}$. Then, we let M_2 be obtained by deleting each $x \in A$ and $y \in B$ and renaming each x_a and y_b as x and y respectively. Let $C_1 = A \cup \{a, p\}$ and $C_2 = B \cup \{b, q\}$.

Note that C_1 is a circuit of M_1 and hence also in M_2 . Moreover, since each of the elements of B has been “lifted” towards b , the set C_1 is also a hyperplane of M_2 .

Note that, with C_1 a circuit-hyperplane, it suffices to show that C_1 is the only cyclic flat of M_2 that contains p but not q . Suppose that $F \neq C_1$ is a cyclic flat of M_2 that contains p but not q . Thus $F \cap C_2 = F \cap (E(M_2) - C_1) \neq \emptyset$. Since F is cyclic and C_2 is a cocircuit, $|C_2 \cap F| \geq 2$ by orthogonality. Since $q \notin F$, the flat F contains an element $y \in B$. Since each element in B is freer than b in M_2 , we have $b \in F$. Similarly $a \in F$. So $F - \{a, b\}$ is a union of cycles in $M_2/\{a, b\}$. However, $M_2/\{a, b\} = M_1/\{a, b\}$. In M_1 , we have a freely placed in the flat $E(M) \cup \{p\}$ and $\{a, p\}$ a series-pair, and hence p is freely placed in $M_1/\{a, b\}$. However, $p \in F - \{a, b\}$, and hence $F - \{a, b\}$ contains a basis of B' of M . Thus $B' \cup \{a, b\} \subseteq F$ is a basis of M_2 , contrary to the fact that q is not contained in the flat F . \square

The following result captures the difference between the matroids $\mathcal{X}(N, p, q; M, A, B)$ and $\mathcal{X}(N_{p \rightarrow q}, p, q; M, A, B)$. It will let us show that when we delete an element in $A \cup \{a\}$ or contract an element in $B \cup \{b\}$ we will get the same minor, Figure 4.2.

Lemma 4.2.8. *Let $\{p, q\}$ be an incomparable pair in N . Let X be a set of elements in $\mathcal{X}(N, p, q; M, A, B)$. If $\mathcal{X}(N, p, q; M, A, B)$ and $\mathcal{X}(N_{p \rightarrow q}, p, q; M, A, B)$ differ in rank on X , then $X \cap (A \cup B \cup \{a, b\}) = A \cup \{a\}$.*

Proof. Let $\mathcal{X} = \mathcal{X}_2(N, p, q; M, A, B)$ and $\mathcal{X}' = \mathcal{X}_2(N_{p \rightarrow q}, p, q; M, A, B)$. Assume that \mathcal{X} and \mathcal{X}' differ in rank on X . As $N_{p \rightarrow q}$ is freer than N , we have that \mathcal{X}' is freer than \mathcal{X} and, hence, $r_{\mathcal{X}'}(X) > r_{\mathcal{X}}(X)$. Let $S_1 = E(N)$, $S_2 = A \cup B \cup \{a, b\}$, $X_1 = X \cap S_1$, $X_2 = X \cap S_2$, and $L = \text{cl}(\{p, q\})$.

For \mathcal{X} and \mathcal{X}' to differ in rank on X it must be the case that N and $N_{p \rightarrow q}$ to differ in rank on $X_1 \cup \{p\}$. Thus X_1 spans p but not q in N .

Note that $r_{\mathcal{X}}(X) = r_{\mathcal{X}}(X_1) + r_{\mathcal{X}}(X_2) - \square_{\mathcal{X}}(X_1, X_2)$ and $r_{\mathcal{X}'}(X) = r_{\mathcal{X}'}(X_1) + r_{\mathcal{X}'}(X_2) - \square_{\mathcal{X}'}(X_1, X_2)$, so $\square_{\mathcal{X}}(X_1, X_2) > \square_{\mathcal{X}'}(X_1, X_2)$. However, $\square_{\mathcal{X}}(X_1, L) = \square_{\mathcal{X}'}(X_1, L) = 1$ and $\square_{\mathcal{X}}(X_2, L) = \square_{\mathcal{X}'}(X_2, L)$. Hence $\square_{\mathcal{X}}(X_1, X_2) = 1$ and $\square_{\mathcal{X}'}(X_1, X_2) = 0$. So X_2 spans p in \mathcal{X} and, since $\mathcal{X}|(S_2 \cup L) = \mathcal{X}'|(S_2 \cup L)$, X_2 also spans p in \mathcal{X}' . Since $\square_{\mathcal{X}'}(X_1, X_2) = 0$, we have that X does not span q in \mathcal{X}' or in \mathcal{X} .

Now the result follows from Lemma 4.2.7. \square

We are now ready to prove that proper minors of $\mathcal{X}(N, p, q; M, A, B)$ are in \mathcal{M} . We will, in fact, prove the following more general result.

Theorem 4.2.9. *Let \mathcal{M} be a natural class of matroids. If*

- (i) *N is an excluded minor for \mathcal{M} ,*
- (ii) *p and q are an incomparable pair of elements in N such that $N_{p \rightarrow q}$ and $N_{q \rightarrow p}$ are contained in \mathcal{M} ,*
- (iii) *M is a matroid in \mathcal{M} with $E(M) \cap E(N) = \emptyset$, and*
- (iv) *(A, B) is a partition of $E(M)$ into bases,*

then each proper minor of $\mathcal{X}(N, p, q; M, A, B)$ is contained in \mathcal{M} .

Proof. Let $\mathcal{X} = \mathcal{X}(N, p, q; M, A, B)$. If $e \in E(N) - \{p, q\}$, then, by construction, $\mathcal{X} \setminus e = \mathcal{X}(N \setminus e, p, q; M, A, B)$ and $\mathcal{X}/e = \mathcal{X}(N/e, p, q; M, A, B)$. Then, since $N \setminus e$, N/e and M are all contained in \mathcal{M} , the minors $\mathcal{X} \setminus e$ and \mathcal{X}/e are also contained in \mathcal{M} .

Now, for $e \in A \cup \{a\}$ and $f \in B \cup \{b\}$, it follows from Lemma 4.2.8 that $\mathcal{X} \setminus e = \mathcal{X}(N_{p \rightarrow q}, p, q; M, A, B) \setminus e$ and $\mathcal{X}/f = \mathcal{X}(N_{p \rightarrow q}, p, q; M, A, B)/f$. Then, since $N_{p \rightarrow q}$ and M are all contained in \mathcal{M} , the minors $\mathcal{X} \setminus e$ and \mathcal{X}/f are also contained in \mathcal{M} .

Finally, since $\mathcal{X}(N, q, p; M, B, A) = \mathcal{X}(N, p, q; M, A, B)$, it follows that, for $e \in A \cup \{a\}$ and $f \in B \cup \{b\}$, the minors \mathcal{X}/e and $\mathcal{X} \setminus f$ are contained in \mathcal{M} . \square

By Lemma 4.2.6 we have that \mathcal{X} is not in \mathcal{M} , and by Lemma 4.2.9 all the proper minors of \mathcal{X} are in \mathcal{M} . Thus $\mathcal{X} = \mathcal{X}(N, p, q; M, A, B)$ is an excluded minor for \mathcal{M} . This was Theorem 4.2.3.

We can now prove Theorem 4.1.1, which we restate here for convenience.

[Theorem 4.1.1 (Campbell, Geelen [8]). *Let \mathcal{M} and \mathcal{N} be natural classes where $\mathcal{M} \subsetneq \mathcal{N}$. If \mathcal{M} is closed under pinned extensions into 3-separations, then each matroid in \mathcal{M} is a minor of an excluded minor of \mathcal{M} that is also in \mathcal{N} .*

Proof. Let M_0 be a matroid in the class \mathcal{M} . By Lemma 4.2.2, there is an M_0 -constructed matroid M , containing M_0 as a minor, and a partition (A, B) of $E(M)$ into two bases. By Lemma 4.2.1, there is an excluded minor N for \mathcal{M} such that N is contained in \mathcal{N} and such that N contains an incomparable pair $\{p, q\}$ of elements where $N_{p \rightarrow q}$ and $N_{q \rightarrow p}$ are both in \mathcal{M} .

Let $\mathcal{X} = \mathcal{X}(N, p, q; M, A, B)$. Note \mathcal{X} contains M as the minor $\mathcal{X}/\{a, b\} \setminus E(N)$. By Theorem 4.2.3, \mathcal{X} is an excluded minor for \mathcal{M} . Moreover, since M and N are both contained in the natural class \mathcal{N} , the matroid \mathcal{X} is also contained in \mathcal{N} . \square

Chapter 5

Computational Complexity

We now consider the feasibility of practical approaches. Can we have an actual technique to determine whether a given matroid (E, \mathcal{I}) is or is not real-representable? If so, how complicated must this technique be?

We will see that we can in fact construct an algorithm to determine real-representability; real-representability is decidable. Indeed, it is well established that, for any field \mathbb{F} , the \mathbb{F} -representability of matroids is decidable if and only if the \mathbb{F} -solvability of systems of Diophantine equations is decidable. However, we will see that even to demonstrate whether or not a matroid (E, \mathcal{I}) on n elements is real-representable requires us to review the independence of an exponential number of sets — still less than the number we may need to consider to determine real-representability in the first place.

[**Theorem 5.3.2** (folklore)]. *Real-representability is not polynomially certifiable.*

[**Theorem 5.5.3** (Ben David, Geelen)]. *Non-real-representability is not polynomially certifiable.*

While complex-representability is a necessary condition that is similarly difficult to demonstrate, prior knowledge of complex-representability does not seem to help:

Theorem 5.0.1 (Campbell). *Real-representability is not polynomially certifiable within the class of complex-representable matroids.*

[**Theorem 5.6.2** (Campbell)]. *Non-real-representability is not polynomially certifiable within the class of complex-representable matroids.*

Instead of Theorem 5.0.1 above, we will prove a more general result that holds for natural classes closed under pinned extensions into 3-separations, see Theorem 5.6.1.

5.1 Decidability; solvability of systems of integer polynomial equations

In Section 2.1, we saw that we can encode elementary algebraic relations between values encoded as points on a line. By introducing intermediate calculations, we could plausibly encode algebraic relationships that are expressible as a set of integer polynomial equations. However, as there may be values that coincide, there is a matroid that encodes each possible collection of coincidences. This type of construction was first discussed in a matroid context by MacLane [24], then White [44, Section 1.7] and Sturmfels [39]. While each author was interested in specific applications, the following holds in general.

Theorem 5.1.1 (folklore). *There is an algorithm to encode a finite set of integer polynomials $S \subseteq \mathbb{Z}[x_1, \dots, x_n]$ as a finite set \mathcal{M}_S of matroids such that for each infinite field \mathbb{F} the polynomials in S have a common root in \mathbb{F} if and only if at least one of the matroids in \mathcal{M}_S is \mathbb{F} -representable.*

Conversely, we can reduce the problem of representing a matroid over a field to solving a integer polynomial system over that field:

Theorem 5.1.2 (Folklore). *There is an algorithm that converts a given matroid M into a finite set of integer polynomials $S_M \subseteq \mathbb{Z}[X]$, such that M is representable over a field \mathbb{F} if and only if the polynomials in S_M have a common root in \mathbb{F} .*

Proof. Assign a variable to each entry of an $r(M) \times E(M)$ -matrix. For this matrix to represent M , each $r(M) \times r(M)$ -submatrix must be invertible if and only if its columns corresponds to the elements of a basis of M . We relate whether each submatrix as being either invertible or singular by the determinant either having a multiplicative inverse or equalling zero respectively. For a field \mathbb{F} , an \mathbb{F} -representation corresponds to a solution of this integer polynomial system over \mathbb{F} . \square

Integer polynomial systems are known to be decidable over a given algebraically closed field by the effective Nullstellensatz [21], and over the reals using quantifier elimination [40]. So, by Theorem 5.1.2, we can also decide representability over these fields.

5.2 Certification

We still do not have a sense of how complicated it is to determine whether or not a given property, P , holds for a pair (E, \mathcal{I}) . Even if the answer is already known, to demonstrate

or prove whether or not P holds will require using some information of (E, \mathcal{I}) — though less than that to determine whether or not P holds. While this is the perspective we will take, it is equivalent to considering non-deterministic algorithms with an “oracle”.

Recall that \mathcal{I} is a collection of subsets of E . An *independence evaluation* with the *independence oracle* of a subset $S \subseteq E$ will declare whether or not S is in \mathcal{I} . For a pair (E, \mathcal{I}) that satisfies a property P , we consider having a “Claimant” that knows everything about (E, \mathcal{I}) that wishes to prove to an “Adjudicator” that (E, \mathcal{I}) satisfies P with few petitions to the independence oracle [14]. For example, to demonstrate that (E, \mathcal{I}) is indeed (E, \mathcal{I}) the Claimant must show the Adjudicator an independence evaluation of each the $2^{|E|}$ subsets of E .

Let (E, \mathcal{I}) satisfy property P . A collection C of subsets of E is said to *certify that (E, \mathcal{I}) satisfies P* when every possible (E, \mathcal{I}') that has the same independent sets in C as (E, \mathcal{I}) also satisfies P . We say that we can *polynomially certify P* when there exists a polynomial f such that, for each (E, \mathcal{I}) we consider that satisfies P , there is a collection of size at most $f(|E|)$ that certifies P for (E, \mathcal{I}) .

5.2.1 Certifying minors

Remark 5.2.1. *For a given matroid N , we only require a constant number of independence evaluations to certify that a matroid $M = (E, \mathcal{I})$ contains N as the minor.*

Specifically, if $N = M/C \setminus D$, then the collection of $2^{|N|}$ sets $\{S \cup C : S \subseteq E - (C \cup D)\}$ certifies that N is a minor of M .

5.2.2 Not polynomially certifiable

To prove that a property P is not polynomially certifiable in \mathcal{M} , we need to construct an infinite family of matroids $\{M_n\}_{n \in I}$ in \mathcal{M} that satisfy P , yet where each M_n differs on only a few independent sets from exponentially (in $|M_n|$) many matroids in \mathcal{M} that do not satisfy P . To show that M_n is not one of these exponentially many matroids not satisfying P will require exponentially many probes.

Our global algebraic structures from Section 2.2 will play a key role in constructing such a family $\{M_n\}_{n \in I}$.

5.3 Real-representability is not polynomially certifiable

Seymour [36] proved that even to demonstrate whether or not a matroid (E, \mathcal{I}) on n elements is $\text{GF}(2)$ -representable may require evaluating the independence of an exponential number of sets. He considers the rank- n binary spike M_n represented by the binary matrix $[I_n | J_n - I_n]$, where I_n and J_n are the $n \times n$ identity and all ones matrices, respectively. This matroid has 2^{n-1} dependent transversals and relaxing any one of these circuit-hyperplanes to a basis yields a non-binary matroid. Thus a certificate that M_n is binary will contain each of these 2^{n-1} dependent transversals. Folklore has it that this result extends to any field, and this technique underlies each of the results from this chapter.

We will explicitly consider a construction for the \mathbb{R} -analogue of Seymour's result.

Theorem 5.3.1 (folklore). *Real-representability is not polynomially certifiable.*

However, as this construction is representable over all infinite fields, we will actually prove the following strengthening. Note that the Adjudicator is not aware that the matroids we are considering are representable over all infinite fields.

Theorem 5.3.2. *For matroids representable over all infinite fields, representability is not polynomially certifiable.*

This sharply contrasts with the class of matroids representable over all fields, known as *regular matroids*. Seymour proved a decomposition theorem for matroids representable over all fields that gives us the following.

Theorem 5.3.3 (Seymour [37]). *Representable-over-all-fields is polynomially certifiable.*

To prove Theorem 5.3.2, we use a spike to encode a sum $\sum_{i=1}^n \alpha_i = 0$ (see section 2.2.1). This sum has exponentially many complementary partial sums that must be confirmed to either both be zero or non-zero. One can check that a similar proof using swirls and products also works.

Proof of Theorem 5.3.2. For an integer $n \geq 4$, consider a representable spike M with a pair of disjoint dependent transversals. By putting M in standard form with respect to

one of these transversals, we may assume that M has legs $[0]_1, [\alpha_1]_1, \dots, [0]_n, [\alpha_n]_n$ for some non-zero $\alpha_1, \dots, \alpha_n$ that sum to zero. That is to say, M has representation

$$\left(\begin{array}{ccccc|ccccc} [0]_1 & [0]_2 & \dots & [0]_{n-1} & [0]_n & [\alpha_1]_1 & [\alpha_2]_2 & \dots & [\alpha_{n-1}]_{n-1} & [\alpha_n]_n \\ \hline 1 & 0 & \dots & 0 & -1 & 1 & 0 & \dots & 0 & -1 \\ 0 & 1 & \dots & 0 & -1 & 0 & 1 & \dots & 0 & -1 \\ \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \dots & \vdots & \vdots \\ 0 & 0 & \dots & 1 & -1 & 0 & 0 & 0 & 1 & -1 \\ \hline 0 & 0 & \dots & 0 & 0 & \alpha_1 & \alpha_2 & \dots & \alpha_{n-1} & \alpha_n \end{array} \right).$$

Consider each partition $S \sqcup T$ of $1, \dots, n$ with $|S|, |T| \geq 2$. As $\sum_{i=1}^n \alpha_i = 0$, the partial sum $\sum_{i \in S} \alpha_i$ is zero if and only if the partial sum $\sum_{i \in T} \alpha_i$ is zero. Geometrically, this means that complementary transversals of M are either both bases or both circuits. However, recall that as long as there are no dependent transversals that only differ on a single leg, we still have a spike, but not necessarily a representable one; see [34, Proposition 1.5.17].

Let M_n be the representable spike with legs $[0]_1, [\alpha_1]_1, \dots, [0]_n, [\alpha_n]_n$ for some non-zero $\alpha_1, \dots, \alpha_n$ in any field, that minimally sum to zero. To certify that M_n is representable, we need to show that there is no independent transversal whose complement is dependent. Thus a certificate of representability contains each of the $2^n - 2 - 2n$ independent transversals that differ on at least two elements from both $\{[0]_1, \dots, [0]_n\}$ and $\{[\alpha_1], \dots, [\alpha_n]\}$.

In any infinite field, we can choose such $\alpha_1, \dots, \alpha_n$ for any integer $n \geq 4$. □

5.4 Certifying non-representability for finite fields

Certifying non-representability shows promise at first. Tutte proved that a matroid is non-binary precisely when it contains a 4-point line as a minor [41]. By Remark 5.2.1, the (non-)independence of only a constant number of sets is required to demonstrate this minor (eight in this case) and thus non-GF(2)-representability.

Similarly, as a result of Rota's Conjecture that representability over each finite field has finitely many excluded minors (a proof of which was announced by Geelen, Gerards, and Whittle [13]):

Theorem 5.4.1. *Let \mathbb{F} be a finite field. Non- \mathbb{F} -representability can be certified with a constant number of probes of the collection of independence sets.*

5.5 Non-real-representability is not polynomially certifiable

This section is based on yet unpublished research in collaboration with Shalev Ben David and Jim Geelen.

Unfortunately, we will see that certifying non-representability is not as tractable over infinite fields.

Theorem 5.5.1 (Ben David, Campbell, Geelen). *For any infinite field \mathbb{F} , non- \mathbb{F} -representability is not polynomially certifiable.*

We will prove the following strengthening of the previous theorem.

Theorem 5.5.2 (Ben David, Campbell, Geelen). *For non-representable matroids, non-[representable-over-all-infinite-fields] is not polynomially certifiable.*

That is to say, there are non-representable matroids that the Claimant cannot polynomially certify to be non- \mathbb{F} -representable for their choice of infinite field \mathbb{F} . Note that this also has the following corollary.

Corollary 5.5.3 (Ben David, Geelen). *Non-representability is not polynomially certifiable.*

To prove Theorem 5.5.2, we build a family of matroids that each of which encode a set of inconsistent algebraic relations. Specifically, the n^{th} matroid M_n'' in this family would encode that $x^m \notin P_m$ where $m = 2^n$ and $P_m = \{\prod_{i=1}^m z_i : z_i \in \{x, y\}\}$. Showing that M_n'' is non-representable amounts to verifying that $x^m \neq \prod_{i=1}^m z_i$ for each $\prod_{i=1}^m z_i \in P_m$.

We first consider a family of swirl-like matroids where we artificially impose the condition that $z_i \in \{x, y\}$ in representations.

5.5.1 A restricted representation problem

For $m \geq 1$, let Ω_m be a rank- $(m + 1)$ swirl-like matroid with no dependent transversals, where the zero-th edge has one element while, for $i \in \{1, 2, \dots, m\}$, the i^{th} edge has two elements. When the swirl-like matroid Ω_m is clear, denote its lines L_0, \dots, L_m and its joints $b_i \in L_{i-1} \cap L_i$ for $i \in \mathbb{Z}_{m+1}$. For $i \in \{1, 2, \dots, m\}$, label the two non-joint elements of L_i by $[x_i]_i$ and $[y_i]_i$, and label the non-vertex element of $L_{m+1} = L_0$ by $[x_0^{-m}]_0$.

Recall α, β are k -algebraically independent in a field \mathbb{F} when, for any non-zero polynomial p over \mathbb{F} of degree at most k , we have $p(\alpha, \beta) \neq 0$.

Recall from Section 2.0.3, that a representation f that corresponds to this encoding is one where, for every $i \in \mathbb{Z}_{m+1}$, $f([\sigma(x_i, y_i)]_i) = f(b_{i-1}) - \sigma(x_i, y_i)f(b_i)$ for every rational function σ . We consider representations where there exist distinct, non-zero α, β in a field \mathbb{F} such that, for all $i \in \mathbb{Z}_{m+1}$, we have evaluated $\{x_i, y_i\}$ as $\{\alpha, \beta\}$, but not necessarily respecting this order. We see that there are no such representations and that a certificate of this non-representability contains all transversals and thus has size at least 2^m .

Lemma 5.5.4. *Let $m \geq 1$. Let Ω_m be the rank- $(m+1)$ swirl-like matroid with no dependent transversals given by the matrix*

$$A(X) = \left(\begin{array}{cccc|cccc|c} b_0 & b_1 & b_2 & \dots & b_m & [x_1]_1 & [y_1]_1 & [x_2]_2 & [y_2]_2 & \dots & [x_m]_m & [y_m]_m & [x_0^{-m}]_0 \\ 1 & 0 & 0 & \dots & 0 & 1 & 1 & 0 & 0 & \dots & 0 & 0 & -x_0^{-m} \\ 0 & 1 & 0 & \dots & 0 & -x_1 & -y_1 & 1 & 1 & \dots & 0 & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 & 0 & -x_2 & -y_2 & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 & 0 \\ 0 & 0 & 0 & \dots & 1 & 0 & 0 & 0 & 0 & \dots & -x_m & -y_m & 1 \end{array} \right)$$

over the field of fractions of $\mathbb{Z}[X]$.

Let \mathbb{F} be a field and α, β be distinct, non-zero values in \mathbb{F} . There is no representation of Ω_m of the form $A(\overline{X})$ for $\overline{X} \subseteq \mathbb{F}$ where $\{x_i, y_i\} = \{\alpha, \beta\}$ for all $i \in \{1, \dots, m, 0\}$. Furthermore, a certificate that no such representation exists contains all transversals of Ω_m and thus has size at least 2^m .

Proof. Suppose that Ω_m has a representation A according to the evaluation $\{x_i, y_i\} = \{\alpha, \beta\}$ for all $i \in \{1, \dots, m, 0\}$. For all $i \in \{1, \dots, m, 0\}$, we can interchange the labelling of each x_i and y_i so we lose no generality in assuming that $x_i = \alpha$ and $y_i = \beta$. Thus

$$A = \left(\begin{array}{cccc|cccc|c} b_0 & b_1 & b_2 & \dots & b_m & [x_1]_1 & [y_1]_1 & [x_2]_2 & [y_2]_2 & \dots & [x_m]_m & [y_m]_m & [x_0^{-m}]_0 \\ 1 & 0 & 0 & \dots & 0 & 1 & 1 & 0 & 0 & \dots & 0 & 0 & -\alpha^{-m} \\ 0 & 1 & 0 & \dots & 0 & -\alpha & -\beta & 1 & 1 & \dots & 0 & 0 & 0 \\ 0 & 0 & 1 & \dots & 0 & 0 & 0 & -\alpha & -\beta & \dots & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 & 0 \\ 0 & 0 & 0 & \dots & 1 & 0 & 0 & 0 & 0 & \dots & -\alpha & -\beta & 1 \end{array} \right).$$

However, the submatrix corresponding to the transversal $T = \{[x_1]_1, \dots, [x_m]_m, [x_0^{-m}]_0\}$ is singular, while Ω_n had no dependent transversals by construction. Thus we can have no such representation f .

Now suppose that we have a certificate $C \subseteq 2^{E(\Omega_m)}$ that there is no representation given by an evaluation $\{x_i, y_i\} = \{\alpha, \beta\}$ for all $i \in \{1, \dots, m, 0\}$. Suppose that C does not contain some transversal T . By interchanging x_i and y_i where necessary, we may assume that $T = \{[x_1]_1, \dots, [x_m]_m, [x_0^{-m}]_0\}$ without losing generality. Let Ω_T be the swirl-like matroid on $E(\Omega_m)$ with the same lines as Ω_m given by taking $(m+1)$ -algebraically independent α and β in \mathbb{F} and assigning x_i to α and y_i to β to get an \mathbb{F} -encoding. This gives us an \mathbb{F} -representation up to row equivalence. In particular, Ω_T has a representation with the matrix A above for α and β that are $(m+1)$ -algebraically independent. Recall that the matrix

$$\begin{pmatrix} 1 & 0 & 0 & \dots & 0 & -\gamma_0 \\ -\gamma_1 & 1 & 0 & \dots & 0 & 0 \\ 0 & -\gamma_2 & 1 & \ddots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & 1 & 0 \\ 0 & 0 & 0 & \dots & -\gamma_m & 1 \end{pmatrix}. \quad (5.1)$$

is singular precisely when its determinant $1 - \prod_{i=1}^m \gamma_i$ is zero. Any transversal S of Ω_T that is not T will contain an element labelled $[y_i]_i$ for some $i \in \{1, 2, \dots, m\}$. Thus the determinant of the submatrix corresponding to $S \neq T$ has a positive degree for β and is therefore not the zero polynomial. As α, β are $(m+1)$ -algebraically independent by assumption, this determinant is non-zero and thus S is an independent transversal. Thus Ω_m and Ω_T agree on the independence of every transversal except for T . As they have the same lines, the independence of all non-transversal sets is also the same for Ω_m and Ω_T . Thus the certificate $C \subseteq 2^{E(\Omega_m)}$ must contain all 2^m transversals T . \square

5.5.2 Enforcing this restricted problem

We will now modify each swirl-like matroid Ω_m from the previous section so as to enforce the constraint that $\{x_i, y_i\}$ is a constant set. To do this, we will make use of the following lemma.

Lemma 5.5.5. *For distinct $\alpha, \beta \in \mathbb{F}$, let $s = \alpha + \beta$ and $t = \alpha\beta - 1$. If $x + y = s$ and $xy - 1 = t$, then $\{x, y\} = \{\alpha, \beta\}$.*

Proof. The quadratic in z ,

$$(z - x)(z - y) = z^2 - sz + (t + 1) = (z - \alpha)(z - \beta)$$

only has two solutions as there is a division algorithm for $\mathbb{F}[z]$. \square

For a fixed $m \geq 1$, we construct a rank- $(m + 2)$ swirl-like matroid Ω'_m from the rank- m swirl-like matroid Ω_m by adding a new line L_{-1} and corresponding joints between L_m and L_0 and adding new elements to existing lines. This new line L_{-1} will contain three non-joint points labelled $[1]_{-1}$, $[s^{-1}]_{-1}$, $[t^{-1}]_{-1}$. For each $i \in \{1, 2, \dots, m\}$, the line L_i will still contain $[x_i]_i, [y_i]_i$ and new non-joints labelled $[1]_i$, $[x_i + y_i]_i$, $[x_i y_i - 1]_i$, and $[x_i y_i]_i$. The line L_0 will still contain $[x_0^{-m}]_0$, as well as new non-joints labelled $[1]_0$, $[x_0 + y_0]_0$, $[x_0 y_0 - 1]_0$, $[x_0 y_0]_0$, but also $[y_0]_0$ and $[x_0]_0, [x_0^2]_0, \dots, [x_0^m]_0$. Counting the joints $b_{-1}, b_0, b_1, \dots, b_m$, note that Ω'_m will have $8m + 11$ elements.

We can now define Ω'_m as the unique rank- $(m + 2)$ swirl-like matroid with lines as described above whose only dependent transversals are $I = \{[1]_{-1}, [1]_0, [1]_1, \dots, [1]_m\}$, and $S_i = (I - \{[1]_{-1}, [1]_i\}) \cup \{[s^{-1}]_{-1}, [x_i + y_i]_i\}$ and $T_i = (I - \{[1]_{-1}, [1]_i\}) \cup \{[t^{-1}]_{-1}, [x_i y_i - 1]_i\}$ for each $i \in \{0, 1, \dots, m\}$.

As the dependent transversal I contains all elements labelled $[1]_i$, this encoding is done *with respect to* I . The dependent transversals S_i and T_i will allow us to apply the algebraic Lemma 5.5.5 to $\{x_i, y_i\}$ to enforce the restricted evaluations from the previous section. By reducing to Lemma 5.5.4, we will prove the following.

Proposition 5.5.6. *Let $m \geq 1$. For any field \mathbb{F} , there is no \mathbb{F} -representation of Ω'_m encoding its labels for some evaluation of $\{s, t\} \cup \{x_i, y_i\}_{i=0}^m$. Furthermore, a certificate of this non-representability contains all transversals and thus has size at least 2^m .*

Proof. Suppose that for some field \mathbb{F} we have a representation f of Ω'_m . Recall from Section 2.2.1, that the restriction of f to a transversal $\{[\gamma_1]_1, \dots, [\gamma_m]_m, [\gamma_{-1}]_{-1}, [\gamma_0]_0\}$ is projectively equivalent to the matrix

$$\begin{pmatrix} [\gamma_1]_1 & [\gamma_2]_2 & [\gamma_3]_3 & \dots & [\gamma_{-1}]_{-1} & [\gamma_0]_0 \\ 1 & 0 & 0 & \dots & 0 & -\gamma_0 \\ -\gamma_1 & 1 & 0 & \dots & 0 & 0 \\ 0 & -\gamma_2 & 1 & \ddots & 0 & 0 \\ \vdots & \vdots & \ddots & \ddots & \ddots & \vdots \\ 0 & 0 & 0 & \ddots & 1 & 0 \\ 0 & 0 & 0 & \dots & -\gamma_{-1} & 1 \end{pmatrix} \quad (5.2)$$

which has determinant $1 - \prod_{i=1}^n \gamma_i$. Thus for each $i \in \{0, 1, \dots, m\}$, the dependent transversals S_i and T_i indicate that $x_i + y_i = s$ and $x_i y_i - 1 = t$. By Lemma 5.5.5, this implies that for some α, β in \mathbb{F} we have $\{x_i, y_i\} = \{\alpha, \beta\}$ for all $i \in \{0, 1, \dots, m\}$. Consider $[1]_{-1}$ in the dependent transversal I we have encoded with respect to. Transversals through $[1]_{-1}$ have $\gamma_{-1} = 1$ in the matrix 5.2 above. Thus contracting $[1]_{-1}$ identifies b_{-1} with b_0 and preserves the labelled encoding of points on the lines L_0, \dots, L_m in the new \mathbb{F} -representation. In particular, this representation contradicts the first part of Lemma 5.5.4.

Now suppose that we have a certificate $C \subseteq 2^{E(\Omega'_m)}$ that there is no representation of Ω'_m that respects its labelling. Suppose that C does not contain $T \cup \{[1]_{-1}\}$ for every transversal T of Ω_m . By Lemma 5.5.4, we can take $(3m + 2)$ -algebraically independent α and β in a field \mathbb{F} such that we have a representation that is consistent with the encoding where $\{x_i, y_i\} = \{\alpha, \beta\}$ for $i \in \{0, 1, \dots, m\}$. Thus $x_i + y_i = \alpha + \beta$ and $x_i y_i - 1 = \alpha\beta - 1$ for $i \in \{0, 1, \dots, m\}$. By assigning s and t to $\alpha + \beta$ and $\alpha\beta - 1$, respectively, we have that S_i and T_i are dependent transversals. For every other transversal besides I , the determinant of its corresponding matrix, (5.2), is a non-zero polynomial expression in α and β of degree at most $(3m + 2)$. Thus by choice of α and β , this expression is non-zero for independent transversals. Thus we have a representation of Ω'_m that respects its labelling. Thus the certificate $C \subseteq 2^{E(\Omega_m)}$ must contain all 2^m transversals arising from transversals of Ω_m . \square

This proposition immediately has the following corollary.

Corollary 5.5.7. *Non-representability over $\mathbb{Q}(X)$ subject to a labelled encoding is not polynomially certifiable.*

However, we would like to prove that non-representability is not polynomially certifiable regardless of whether we are subject to a labelled encoding. We do this by by “gluing” von Staudt matroids to impose the necessary algebraic relations on each line, see Section 2.1. Recall that a sufficient condition to define the *proper amalgam of M and M' across F* is for the set F are “gluing” across to be a modular flat of M' (see Section 2.3).

Proof of Theorem 5.5.2. Let $n \geq 1$ We will construct M_n by starting with Ω'_n and taking a sequence of proper amalgams with the von Staudt matroids \mathcal{O}^+ , \mathcal{O}^* , and \mathcal{O}' to encode the algebraic relations on each line L_i of Ω'_n . We proceed as follows. For each edge $i \in \{0, 1, 2, \dots, n\}$, glue \mathcal{O}^+ to $b_i = [0]_i, b_{i+1} = [-\infty]_i, [x_i]_i, [y_i]_i, [x_i + y_i]_i$, appropriately; glue \mathcal{O}^* to $b_i = [0]_i, b_{i+1} = [-\infty]_i, [x_i]_i, [y_i]_i, [1]_i, [x_i y_i]_i$, appropriately; and finally glue \mathcal{O}^+ to $b_i = [0]_i, b_{i+1} = [-\infty]_i, [1 - x_i y_i]_i, [x_i y_i]_i, [1]_i$, appropriately. Along the line L_0 , we additionally glue a copy of \mathcal{O}' to $b_0 = [0]_0, b_1 = [-\infty]_0, [x_0]_0, [1]_0, [x_0^2]_0$; for each $j \in \{3, \dots, n\}$, a copy of \mathcal{O}^* to $b_0 = [0]_0, b_1 = [-\infty]_0, [x_0^{j-1}]_0, [x_0]_0, [1]_0, [x_0^j]_0$; and finally a copy

of \mathcal{O}' to $b_0 = [0]_0, b_1 = [-\infty]_0, [x_0^n]_0, [x_i^{-n}]_0$. Call each line L_i of M_n together with the matroids attached to M_n at L_i a *petal* of M_n , and label it by P_i .

Suppose that we have a representation f for M_n over a field \mathbb{F} . As projective equivalence preserves the matroid, we may assume that f respects the labels of the dependent transversal $I = \{[1]_{-1}, [1]_0, [1]_1, \dots, [1]_n\}$, (see Section 2.2.2). That is to say, $f([1]_i) = f([0]_i) - f([-\infty]_i)$ for all $i \in \mathbb{Z}_{n+2}$. Note for all $i \in \{0, \dots, n\}$, we have the $[x_i]_i, [y_i]_i, [0]_i$, and $[-\infty]_i$ on the same line but with no parallel pairs. Thus $f([x_i]_i) = f([0]_i) - x_i f([-\infty]_i)$ and $f([y_i]_i) = f([0]_i) - y_i f([-\infty]_i)$ for some distinct non-zero x_i, y_i in \mathbb{F} . By this choice, we have that f encodes the labels $[x_i]_i$ and $[y_i]_i$ for all $i \in \{0, \dots, n\}$. By the von Staudt Lemmas 2.1.1, 2.1.2, 2.1.3 applied in order to the matroids $\mathcal{O}^+, \mathcal{O}^*$, and \mathcal{O}' we glued to Ω'_n , we have that f respects all the labels on L_0, \dots, L_n . Finally, by taking $s = x_0 + y_0$ and $t = x_0 y_0 - 1$, we have that f is subject to encoding Ω'_n according to its labelling. Thus, by Proposition 5.5.6, no representation of M_n exists.

Suppose that we have a certificate C that M_n is not representable over some infinite field \mathbb{F} . We now will construct a certificate C_Ω that Ω'_n is not representable subject to its labelled encoding. For each set $A \in C$ and each $i \in \mathbb{Z}_{n+2}$, let A_i be the set of points in L_i that are spanned by $A \cap P_i$. Let $A_\Omega = \bigcup_{i \in \mathbb{Z}_{n+2}} A_i$. Recall M_n is formed by taking the amalgam of Ω'_n and a petal P_i at each line L_i of M_n and that $E(\Omega'_n) = \bigcup_{i \in \mathbb{Z}_{n+2}} L_i$. Thus, the collection $\{A \cap P_i\}_{i \in \mathbb{Z}_{n+2}}$ of the restrictions of A to each petal and the set A_Ω of points spanned by these restrictions are enough to determine the independence of A . Specifically, A is independent in M_n precisely when A_Ω is independent in Ω'_n and $A \cap P_i$ is independent in P_i for each $i \in \mathbb{Z}_{n+2}$. Let $C_\Omega = \{A_\Omega : A \in C\}$ and note $|C_\Omega| \leq |C|$.

Suppose that C_Ω is not a certificate that Ω'_n is not representable subject to its labelled encoding. Then there is a swirl-like matroid $\overline{\Omega'_C}$ with the same labelling as Ω'_n and an \mathbb{F} -representation f that respects this labelling but that may differ on the independence of transversals not in C_Ω . As f respects the labelling of each line L_i , we can extend f to a representation of each petal P_i by the easy direction of the von Staudt Lemmas 2.1.1, 2.1.2, 2.1.3. This extended \mathbb{F} -representation defines a matroid $\overline{M_C}$ with $E(\overline{M_C}) = E(M_n)$. Since $\overline{M_C}$ and M_n have the $\overline{\text{set}}$ of same petals $\{P_i\}_{i \in \mathbb{Z}_{n+2}}$, for $A \in C$, the restriction of A to each petal is the same in $\overline{M_C}$ and M_n . As the on the independence of each set $A \in C$ is determined by these restrictions and the set A_Ω in C_Ω , we have that $\overline{M_C}$ and M_n agree on the independence of each $A \in C$. Thus contradicts the assumption that C is a certificate of non- \mathbb{F} -representability. So C_Ω is indeed a certificate that Ω'_n is not representable subject to its labelled encoding. By Proposition 5.5.6, C_Ω has size at least 2^n and hence C has size at least 2^n . \square

5.6 Assuming complex-representability

By Theorem 5.3.2, we have that complex-representability is not polynomially certifiable. What if we assume complex-representability is known to the Adjudicator when the Claimant is trying to certify whether or not a matroid is real-representable?

Theorem 5.5.2 already tells us that non-complex-representability is not polynomially-certify when non-real-representability is assumed.

5.6.1 Real-representability is still not polynomially certifiable

Recall that a *pinned extension into a 3-separation* includes a new element that is “pinned” on either side of the 3-separation, see Section 2.3.1.

We prove the following generalization of Theorem 5.0.1.

Theorem 5.6.1. *Let \mathcal{M} and \mathcal{N} be classes of matroids closed under isomorphism, adding coloops, deletion, and principal extension. Let \mathcal{M} also be closed under pinned extensions into 3-separations. If $\mathcal{N} - \mathcal{M} \neq \emptyset$, then membership in \mathcal{M} is not polynomially certifiable even within the class \mathcal{N} .*

Notably \mathcal{M} may be a class matroids representable over a fixed infinite field \mathbb{F} and \mathcal{N} may be any “natural” class (see Section 4.1). However, as we do not require that \mathcal{M} and \mathcal{N} be closed under direct sums, \mathcal{M} and \mathcal{N} may also be arbitrary unions (and intersections) of these classes.

Proof of Theorem 5.6.1. Let N be a matroid in $\mathcal{N} - \mathcal{M}$ with $|E(N)|$ minimum and, subject to this, with N freest on ground set $E(N)$. Let p and q be distinct elements of N and let N' denote the matroid obtained from N by freely adding a new element e into the line spanned by $\{p, q\}$. Let $L \supseteq \{p, e, q\}$ be the line spanned by $\{p, q\}$ in N' . Note that e is freer than p in N' , so there is a subset Y_p of $E(N') - \{p\}$ that spans p and not e in N' . We have $N \wedge p$ in \mathcal{M} , as otherwise renaming e as p in $N \wedge p$ contradicts the choice of N . However, $N' \notin \mathcal{M}$ as it contains $N \notin \mathcal{M}$ as a restriction.

For $n \geq 1$, we construct the matroid $M_n \in \mathcal{M}$ as follows. Start with $N \wedge p$ in \mathcal{M} and add n coloops a_1, \dots, a_n and then freely place $a_{n+1}, b_1, \dots, b_n, b_{n+1}$ in the rank- $(n+2)$ flat $(L - \{p\}) \cup \{a_1, \dots, a_n\}$ given by union of the line $L - \{p\}$ and these new coloops. Call this new matroid M_n . Since \mathcal{M} is closed under adding coloops and principal extensions, we have $M_n \in \mathcal{M}$.

Consider $U = \{a_1, \dots, a_{n+1}, b_1, \dots, b_{n+1}\}$ in M_n . Note each element of U was either a coloop or freely placed in the span of the line $L - \{p\}$ and these n coloops. Thus an element of U is contained in a cyclic flat if and only if that cyclic flat contains the flat $(L - \{p\}) \cup U$. This means that we do not change the structure of M_n by interchanging the label of elements in U .

Suppose that we have a certificate $C \subseteq 2^{E(M_n)}$ that M_n is in \mathcal{M} given that $M_n \in \mathcal{N}$. We now show that $\{U \cap T : T \in C\}$ contains all subsets of U of size $n+1$. This will give us that $|C| > \binom{2(n+1)}{n+1} > 2^{n+1}$. For some subset A of U of size $n+1$, suppose that C does not contain any set T with $T \cap U = A$. By interchanging the labels of the elements in U , we may assume that $A = \{a_1, \dots, a_n, a_{n+1}\}$ without losing generality. We will now construct a matroid $M_A \in \mathcal{N} - \mathcal{M}$ that agrees with M_n on the independence of sets in C .

Take $N' \in \mathcal{N} - \mathcal{M}$ as previously described. As before add n coloops a_1, \dots, a_n and then freely place b_1, \dots, b_n, b_{n+1} in the rank- $(n+2)$ flat $L \cup \{a_1, \dots, a_n\}$ given by union of the line L and the new coloops. However, we now freely place a_{n+1} in the rank- $(n+1)$ flat $\{q\} \cup \{a_1, \dots, a_n\}$ formed by q and the new coloops. Finally, we delete q to get M_A .

Since \mathcal{N} is closed under adding coloops and principal extensions, we have $M_A \in \mathcal{N}$. However M_A is not in \mathcal{M} . Note that $(E(N \setminus p), U)$ is a 3-separation of M_A , and that $\square(E(N \setminus p), A) = \square(Y_p, A) = \square(Y_p, U) = 1$, so (Y_p, A) pins p . If M_A was in \mathcal{M} , then as \mathcal{M} is closed under pinned extensions into 3-separations by assumption, we could uniquely “reinsert” p by Lemma 2.3.5. This would give us a matroid in \mathcal{M} that contains $N \in \mathcal{N} - \mathcal{M}$ as a restriction, a contradiction as \mathcal{M} is restriction closed.

Note that for M_n , the element a_{n+1} was freely placed in the flat $L \cup \{a_1, \dots, a_n\} \cup \{b_1, \dots, b_n, b_{n+1}\}$ of N' . However for M_A , the element a_{n+1} was freely placed in $\{q\} \cup \{a_1, \dots, a_n\}$, a hyperplane of the restriction to $L \cup \{a_1, \dots, a_n\} \cup \{b_1, \dots, b_n, b_{n+1}\}$. Thus M_n and M_A only differ on the independence of sets of the form $T = \{a_{n+1}\} \cup S$, where S would span the new coloops a_1, \dots, a_n but none of the elements b_1, \dots, b_n, b_{n+1} . Thus $T \cap U = \{a_1, \dots, a_n, a_{n+1}\} = A$, but we assumed that C contained no such set.

This contradicts the assumption that C is a certificate that M_n is in \mathcal{M} assuming M_n is in \mathcal{N} . So $\{U \cap T : T \in C\}$ contains all subsets of U of size $n+1$, and so the certificate C has size at least 2^n . Since $|E(M_n)| = |E(N)| + 2(n+1)$ for positive integer n , this is not bounded by a polynomial in $|E(M_n)|$. \square

5.6.2 Non-real-representability is still not polynomially certifiable

Theorem 5.6.2 (Campbell). *Non-real-representability is not polynomially certifiable even within the class of complex-representable matroids.*

We can modify the construction that we used to show that non-representability is not polynomially certifiable (Theorem 5.5.2). Recall this construction relied on having to show that $\alpha^n \notin \prod_{i=1}^n \{\alpha, \beta\}$ to certify non-representability. However, the new construction will rely on having to show that $\gamma \notin \prod_{i=1}^n \{\alpha, \beta\}$ where $\gamma^3 = \alpha^{3n}$. Note that this algebraic constraint is impossible for the reals where the cube root is unique, but may be satisfied for complex numbers when $\gamma = \omega\alpha^n$ for a complex 3rd root of unity ω . Specifically, for $m \geq 1$, the line L_0 of Ω''_m will still contain $[1]_0, [x_0 + y_0]_0, [1 - x_0y_0]_0, [x_0y_0]_0$, but will differ from Ω'_m by having $[x_0]_0, [x_0^2]_0, \dots, [x_0^{3m}]_0$ and $[z]_0, [z^2]_0, [z^3]_0 = [x_0^{3m}]_0$, as well as $[z^{-1}]_0$ instead of $[x_0^{-m}]_0$. The matroid M_n is then constructed by gluing von Staudt constructions $\mathcal{O}^+, \mathcal{O}^*$, and \mathcal{O}' appropriately. The proof proceeds essentially the same way as that of Theorem 5.5.2, where z only needs to be evaluated as x_0^m when considering the real representation.

Chapter 6

Orientability

In the previous sections we have seen that characterizing real-representability from first-principles is difficult. This indicates that we need a similarly complicated matroid property to have a non-tautological characterization of real-representability. We saw that even with knowledge of complex-representability, characterizations fail spectacularly. We will shortly define a matroid property, “orientability”, that is also necessary for real-representability. However, we will see that even with complex-representability and orientability, fundamental characterizations of real-representability continue to fail spectacularly.

Given a real-representation of a matroid M , an orientation of M is naturally induced by partitioning each circuit according to the signs in the linear dependency of the corresponding vectors, see Theorem 6.0.3. In this way, we can think of an orientation as a record of the “signs” in each linear dependency. This interpretation only makes sense when the matroid is representable, and there do exist non-representable orientable matroids such as the Vámos matroid [2]. However, a representation and an orientation together may induce a representation over an ordered field. While there are many negative results for representability, this would give the following promising potential characterization proposed by Whittle (private communication, 2017).

Conjecture 6.0.1 (Whittle). *A matroid is real-representable if and only if it is orientable and representable over some field.*

While the forward direction is trivial, the converse direction would be spectacular. This conjecture had already been shown to hold for binary and ternary matroids, with a precise characterization in each of these cases, see [2] and [23], respectively. However, this conjecture does not hold in general.

Theorem 6.0.2 (Campbell, Geelen). *There exist a matroid that is complex-representable and orientable but not representable over the reals.*

In fact, we will see that:

[Theorem 6.2.1 (Campbell)]. *For every finite field \mathbb{F} with $|\mathbb{F}^*| = |\mathbb{F}| - 1$ composite, there is an \mathbb{F} -representable, complex-representable, orientable matroid that is not representable over the reals.*

We will then see that we have the orientable-matroid generalization of the main theorems from Sections 3, 5, and 4. This will either be because we can apply a generalization from that Section to a matroid given by Theorem 6.2.1, or simply because the construction used in the proof is still valid.

6.0.1 Orienting circuits

Recall that a *circuit* in a matroid is a minimal dependent set. Let $M = (E, \mathcal{C})$ be a matroid as described by its set, \mathcal{C} , of circuits. The fundamental conditions that \mathcal{C} must satisfy and which axiomatize circuit descriptions of matroids are as follows.

- (C1) The empty set is not in \mathcal{C} .
- (C2) No proper subset of an element of \mathcal{C} is also in \mathcal{C} .
- (C3) If $C_1, C_2 \in \mathcal{C}$ and there is $e \in C_1 \cap C_2$, then there exists $C_3 \in \mathcal{C}$ such that $C_3 \subseteq (C_1 \cup C_2) - \{e\}$.

The last property, (C3), is known as *circuit elimination property*.

Now we consider why (C3) is true for minimal **linearly** dependent subsets of a set, $\{\mathbf{v}_p\}_{p \in E}$ of vectors over a field \mathbb{F} . Suppose we have minimal linear dependencies $\sum_{p \in C_1} a_p \mathbf{v}_p = 0$ and $\sum_{p \in C_2} b_p \mathbf{v}_p = 0$ for some non-zero coefficients $\{a_p\}_{p \in C_1} \subseteq \mathbb{F}$ and $\{b_p\}_{p \in C_2} \subseteq \mathbb{F}$. If there is $e \in C_1 \cap C_2$, then we can get the linear dependency

$$\begin{aligned} & \frac{1}{a_e} \left(\sum_{p \in C_1} a_p \mathbf{v}_p \right) - \frac{1}{b_e} \left(\sum_{p \in C_2} b_p \mathbf{v}_p \right) = 0 \\ \left(\sum_{p \in C_1 - C_2} \frac{a_p}{a_e} \mathbf{v}_p \right) + \left(\sum_{p \in (C_1 \cap C_2) - \{e\}} \left(\frac{a_p}{a_e} + \frac{b_p}{-b_e} \right) \mathbf{v}_p \right) + \left(\sum_{p \in C_2 - C_1} \frac{b_p}{-b_e} \mathbf{v}_p \right) = 0 \end{aligned} \quad (6.1)$$

which eliminates e . Thus, there is a minimal linear-dependent set of vectors indexed by a subset, C_3 , of $(C_1 \cup C_2) - \{e\}$, the non-zero coefficients in (6.1).

Consider when \mathbb{F} is the reals or another “ordered” field. Then we have a notion of the signs of the non-zero coefficients $\{a_p\}_{p \in C_1}$ and $\{b_p\}_{p \in C_2}$, which induce partitions of C_1 and C_2 . Multiplying these coefficients by $\frac{1}{a_e}$ and $-\frac{1}{b_e}$, respectively, to get $\{\frac{a_p}{a_e}\}_{p \in C_1}$ and $\{-\frac{b_p}{b_e}\}_{p \in C_2}$ does not change these partitions by signs but does ensure that \mathbf{v}_e has coefficients with opposite signs in the new linear dependencies. Now by eliminating \mathbf{v}_e , we find the linear dependency (6.1) indexed by a subset of $(C_1 \cup C_2) - \{e\}$. The signs of this linear dependency must be consistent with the signs of the original linear dependencies. As we will see below, there is a minimal linear dependency whose signs agree with (6.1). It is in this way we may “orient” the circuits of a set $\{\mathbf{v}_p\}_{p \in E}$ of real vectors.

More generally, an *orientation* of the circuits \mathcal{C} of a matroid M is a collection $\mathcal{S} = \{\{C^1, C^2\} : C \in \mathcal{C}\}$, where $\{C^1, C^2\}$ is a partition of C , and such that if C_1 and C_2 are distinct circuits with partitions $\{C'_1, C''_1\}$ and $\{C'_2, C''_2\}$ in \mathcal{S} respectively and with some $e \in C'_1 \cap C''_2$, then there is a circuit C_3 with partition $\{C'_3, C''_3\}$ in \mathcal{S} such that $C'_3 \subseteq (C'_1 \cup C'_2) - \{e\}$ and $C''_3 \subseteq C''_1 \cup C''_2 - \{e\}$. We say that a matroid (E, \mathcal{C}) is *orientable* when there exists an orientation of \mathcal{C} . We call (E, \mathcal{S}) an *oriented matroid*.

Oriented matroids were independently developed by Bland, Folkman, Las Vergnas, and Lawrence, with each making significant contributions. Two joint papers were published in 1978 in the Journal of Combinatorial Theory, Series B, one by Bland and Las Vergnas [2] and the other by Lawrence with Folkman [11].

Theorem 6.0.3 (Bland, Las Vergnas [2]; Folkman, Lawrence [11]). *All real-representable matroids are orientable.*

Proof. Let M be a matroid with real representation $\{\mathbf{v}_p\}_{p \in E}$. For a circuit C of M , let $\{C^1, C^2\}$ be a partition of signs in a linear dependency of $\{\mathbf{v}_p\}_{p \in C}$. Note that this partitioning by signs is unique, as otherwise we could eliminate a vector in the linear dependency and contradict the minimality of C .

Let C_1 and C_2 be distinct circuits in M . Let the signs of the linear dependencies of $\{\mathbf{v}_p\}_{p \in C_1}$ and $\{\mathbf{v}_p\}_{p \in C_2}$ induce the partitions $\{C'_1, C''_1\}$ and $\{C'_2, C''_2\}$ in \mathcal{S} respectively. If $e \in C'_1 \cap C''_2$, then as before consider the linear dependency (6.1) of $\{\mathbf{v}_p\}_{p \in (C_1 \cup C_2) - \{e\}}$ obtained by eliminating e . Note that the partitioning of $(C_1 \cup C_2) - \{e\}$ induced by the signs in this linear dependency satisfies the necessary conditions for an orientation. However, this linear dependency may not be minimal. We now see that there is a minimal linear dependency whose non-zero coefficients agree in sign with the signs in the linear dependency (6.1).

Let $C_3 \subseteq (C_1 \cup C_2) - \{e\}$ be the index set of a linear dependency whose sign-partitioning agrees with the signs of (6.1) and which is minimal subject to this condition. Let \overline{C}_3 be a subset of C_3 that indexes a linear dependency that minimally disagrees in signs with that of C_3 . If $\overline{C}_3 \subsetneq C_3$, then we can eliminate an element that has different signs in the dependencies given by \overline{C}_3 and C_3 and contradict the minimality of \overline{C}_3 . Thus $C_3 \subseteq (C_1 \cup C_2) - \{e\}$ is a minimal linear dependency and whose sign-partitioning agrees with that of C_1 and C_2 in the elimination of e . \square

For a real-representable matroid $M = (E, \mathcal{C})$, consider an orientation \mathcal{S} induced by a real-representation $\{\mathbf{v}_p\}_{p \in E}$ of M . That is to say, a circuit $C \in \mathcal{C}$ has the unique partition $\{C^1, C^2\} \in \mathcal{S}$ where

$$\sum_{p \in C^1} a_p \mathbf{v}_p = \sum_{q \in C^2} b_q \mathbf{v}_q \quad (6.2)$$

with $\{a_p\}_{p \in C^1}$ and $\{b_q\}_{q \in C^2}$ sets of non-zero positive reals. Note that if we scale a vector \mathbf{v}_e in $\{\mathbf{v}_p\}_{p \in E}$ by a negative number, we must *reorient* each circuit, C , that contains e by changing the which part of the partition $\{C^1, C^2\}$ contains e . By row operations and column scaling — reorienting \mathcal{S} appropriately, we may assume that the first entry of each \mathbf{v}_p is one. Interpreting the remaining entries of each \mathbf{v}_p as Euclidean coordinates of a point P_p , the relation (6.2) implies that the convex hull of $\{P_p\}_{p \in C^1}$ intersects the convex hull of $\{P_q\}_{q \in C^2}$. In this way we can interpret orientations as further describing how elements of a matroid are “arranged”. In particular, consider a line L on which we have encoded according to a real-representation f as described in Section 2.0.3. By projective equivalence and reorienting the orientation \mathcal{S} appropriately, we may assume that $f([0]) = \mathbf{e}_1$, $f([\pm\infty]) = \pm\mathbf{e}_2$, and $f([\alpha]) = \mathbf{e}_1 + \alpha\mathbf{e}_2$ for $\alpha \in \mathbb{R}$. For distinct reals $\alpha > \beta > \gamma$, we have $\frac{\alpha-\beta}{\alpha-\gamma}$ and $\frac{\beta-\gamma}{\alpha-\gamma}$ greater than zero and

$$f([\beta]) = \mathbf{e}_1 + \beta\mathbf{e}_2 = \frac{\beta-\gamma}{\alpha-\gamma} (\mathbf{e}_1 + \alpha\mathbf{e}_2) + \frac{\alpha-\beta}{\alpha-\gamma} (\mathbf{e}_1 + \gamma\mathbf{e}_2) = \frac{\beta-\gamma}{\alpha-\gamma} f([\alpha]) + \frac{\alpha-\beta}{\alpha-\gamma} f([\gamma]).$$

Thus we have the orientation $\{\{[\beta]\}, \{[\alpha], [\gamma]\}\}$ in \mathcal{S} . Similarly $f([\beta]) = \mathbf{e}_1 + \beta\mathbf{e}_2 = (\beta-\gamma)f([\infty]) + f([\gamma])$ and $f([\beta]) = \mathbf{e}_1 + \beta\mathbf{e}_2 = f([\alpha]) + (\alpha-\beta)f([-\infty])$, So we have $\{\{[\beta]\}, \{[\infty], [\gamma]\}\}$ or $\{\{[\beta]\}, \{[\alpha], [-\infty]\}\}$ in \mathcal{S} . In this way an orientation \mathcal{S} imposes an ordering of encodings on lines.

6.0.2 Orienting the complements of hyperplanes

A *cocircuit* of a matroid M is a minimal subset G of $E(M)$ that intersects every basis of M . Note that a hyperplane of M is a maximal subset of $E(M)$ that does not contain a basis.

This implies that $G \subset E(M)$ is a cocircuit if and only if $E(M) - G$ is a hyperplane. One can check that the set, \mathcal{C}^* , of cocircuits of a matroid M satisfies analogous properties as the properties (C1), (C2), and (C3) for circuits from the previous section. The cocircuits of M are the circuits of a matroid M^* on the same ground set that is called the *dual* of M .

Recall that we had an “orthogonal vector” for any hyperplane in a represented matroid:

[Remark 2.3.2]. *If M is a representable matroid with representation $\{\mathbf{v}_e\}_{e \in E(M)}$, then for any hyperplane H of M there is a vector \mathbf{w}_H such that $H = \{e \in E(M) : (\mathbf{w}_H)^T \mathbf{v}_e = 0\}$.*

For a matroid M with real-representation $\{\mathbf{v}_e\}_{e \in E(M)}$, we can partition the complement of a hyperplane H according to the sign of each $\mathbf{v}_e \cdot \mathbf{w}_H$ for $e \in (E(M) - H)$, for some choice of orthogonal vector \mathbf{w}_H of H . That is to say, the cocircuit $E(M) - H$ has bipartition

$$\{\{e \in E(M) - H : \mathbf{v}_e \cdot \mathbf{w}_H > 0\}, \{e \in E(M) - H : \mathbf{v}_e \cdot \mathbf{w}_H < 0\}\}.$$

In this way, the set, \mathcal{C}^* , of cocircuits of a matroid M has an orientation \mathcal{S}^* that is analogous to the orientation of the set of circuits.

6.1 Orientability-preserving operations

Knowing that real-representable matroids are orientable (Theorem 6.0.3) gives us a way to start with an oriented matroid. It is known that the class of orientable matroids is closed under isomorphism, minors, direct sums, adding coloops, and principal extension, see [1, Proposition 7.9.1]. That is to say, oriented matroids are a “natural class” as defined in Section 4.1. This will allow us to modify orientable matroids and will be particularly useful when we wish to strengthen the results from the previous sections. However, real-representability is also a “natural class”; we still need techniques to create non-real-representable matroids while preserving orientability.

For this we will use “circuit-hyperplane relaxations” and generalized parallel connection.

6.1.1 Circuit-hyperplane relaxation

If H is both a hyperplane and a circuit in a matroid M , we call H a *circuit-hyperplane* of M . If H is a circuit-hyperplane of a matroid $M = (E, \mathcal{B})$ as described by its bases, it is well known that $(E, \mathcal{B} \cup \{H\})$ will also describe the bases of a matroid. This new matroid,

$M' = (E, \mathcal{B} \cup \{H\})$, is known as the *circuit-hyperplane relaxation* of H in M . In terms of circuits, it can be checked that if the circuit-hyperplane relaxation of H in $M = (E, \mathcal{C})$ yields $M' = (E, \mathcal{C}')$, then $\mathcal{C}' = (\mathcal{C} - \{H\}) \cup \{H \cup \{e\} : e \in (E - H)\}$, see [34, Proposition 1.5.14]. If M has an orientation \mathcal{S} , where the circuit H is oriented $\{H^1, H^2\}$ and the cocircuit $G = E - H$ is oriented $\{G^1, G^2\}$, then we can view relaxation as “perturbing” the points in H^1 towards the G^1 “side” of H and the points in H^2 towards the G^2 “side”. This would give us the orientation

$$\mathcal{S}' = (\mathcal{S} - \{\{H_1, H_2\}\}) \sqcup \{\{H_1, H_2 \cup \{e\}\} : e \in G_1\} \sqcup \{\{H_1 \cup \{e\}, H_2\} : e \in G_2\}$$

for $M' = (E, \mathcal{C}')$, the relaxation of H in M .

In this way, we can prove that

Theorem 6.1.1 (Edmonds, Mandel [10]). *The class of orientable matroids is closed under circuit-hyperplane relaxation.*

6.1.2 Amalgams of orientable matroids

Lemma 6.1.2 (Hochstättler, Nickel [16]). *Let M_1 and M_2 be orientable matroids on ground sets E_1 and E_2 , respectively. Let $L = E_1 \cap E_2$ be a common modular line of M_1 and M_2 and $R = M_1|L = M_2|L$. Let M_1 and M_2 have orientations \mathcal{S}_1 and \mathcal{S}_2 , respectively, such that $\mathcal{S}_1|L = \mathcal{S}_2|L$. Then the general parallel connection $P_R(M_1, M_2)$ has an orientation \mathcal{S} such that $\mathcal{S}|E_1 = \mathcal{S}_1$ and $\mathcal{S}|E_2 = \mathcal{S}_2$.*

We can use this to show the following.

Lemma 6.1.3. *Let M_1 and M_2 have orientations \mathcal{S}_1 and \mathcal{S}_2 and modular lines L_1 and L_2 , respectively. Let $L = L_1 \cap L_2 = E(M_1) \cap E(M_2)$ and have common restriction $R = M_1|L = M_2|L$ of rank 2. If $\mathcal{S}_1|L = \mathcal{S}_2|L$, then there is an orientation \mathcal{S} of $M_1 \oplus_R M_2$ for which $\mathcal{S}|E(M_i) = \mathcal{S}_i$ for $i \in \{1, 2\}$.*

As how Theorem 2.3.3 was used to prove Lemma 2.3.4, we can use Lemma 6.1.2 to prove Lemma 6.1.3 by first appropriately placing the elements of $L_2 - L_1$ on L_1 in M_1 and elements of $L_1 - L_2$ on L_2 in M_2 .

6.2 Representable orientable matroids that are not real-representable

Theorem 6.2.1. *For every finite field \mathbb{F} with $|\mathbb{F}^*| = |\mathbb{F}| - 1$ composite, there is an \mathbb{F} -representable, orientable matroid that is not real-representable.*

In other words, for each prime power $q \geq 5$, if q is not 2^n for some integer n where $2^n - 1$ is a (Marsenne) prime, then there is a $\text{GF}(q)$ -representable matroid that is orientable and complex-representable but not real-representable.

By Moore's classification of finite fields, we may assume that \mathbb{F} is $\text{GF}(q)$, where q is a prime power [30]. As $x^t - 1$ can have at most t roots in $\text{GF}(q)$ for any integer $t \geq 1$, and we can multiply elements of relatively prime orders to get an element with a larger order, we have that $\text{GF}(q)$ has a generator and hence $\text{GF}(q)^* \cong \mathbb{Z}_{q-1}$.

We will prove Theorem 6.2.1 by construction. However the type of construction will depend on whether q is odd or even. If $q \geq 4$ is odd, say $q - 1 = 2k$ with $k \geq 2$, we will construct a swirl that encodes that $\text{ord}(\alpha) = 2k$ for some value α , and while this is possible over $\text{GF}(q)$ and the complex numbers, it is not possible over the reals. This swirl will be orientable as it can be obtained from a real-representable matroid with a sequence of circuit-hyperplane relaxations. If instead $q \geq 4$ is even with $q - 1$ composite, say $q - 1 = st$ for odd integers $s > 1$ and $t > 1$ with s minimal. As in Section 3.4.2, we will construct the amalgam of two real-representable swirls that encodes $\alpha^s = \beta^s$ with $\alpha \neq \beta$. This is not possible over the reals as s is odd, but in $\text{GF}(q)$ or \mathbb{C} we can take α as a $(q - 1)$ -th primitive root of unity and $\beta = \alpha^{t+1}$.

As it involves a more elegant construction, we will first consider the case when q is odd:

Theorem 6.2.2. *For every odd prime power $q > 3$, there is a $\text{GF}(q)$ -representable, orientable matroid that is not real-representable.*

Proof. Suppose $q \geq 4$ is odd. So $q - 1 = 2k$ for some $k \geq 2$.

Consider the $\mathbb{Q}(x)$ -matrix

$$A_q(x) = \begin{pmatrix} [1]_1 & [x]_1 & [1]_2 & [x]_2 & \dots & [1]_{q-1} & [x]_{q-1} & [1]_q & [x]_q \\ 1 & 1 & 0 & 0 & \dots & 0 & 0 & -1 & -x \\ -1 & -x & 1 & 1 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & -1 & -x & \dots & 0 & 0 & 0 & 0 \\ \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \\ 0 & 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & 1 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & \dots & -1 & -x & 1 & 1 \end{pmatrix}.$$

Let W_q be the swirl-like matroid with these same elements, with $[1]_i, [x]_i$ on the i -th line L_i for $i \in \{1, \dots, q\}$ and with the dependent transversals $I = \{[1]_1, \dots, [1]_q\}$ and, for each $i \in \{1, \dots, q\}$, $T_i = (\{[x]_1, \dots, [x]_q\} - \{[x]_i\}) \cup \{[1]_i\}$. Note each T_i has $q - 1 = 2k$ elements of the form $[x]_j$ for some $j \in \{1, \dots, q\}$.

6.2.2.1. W_q is orientable.

Proof. Consider the real-representable swirl-like matroid N_q with representation $A_q(-1)$. By Theorem 6.0.3, we have an orientation of N_q . Note the dependent transversals of N_q are those that contain an even number of elements of the form $[-1]_j$ for some $j \in \{1, \dots, q\}$. In particular, we have the dependent transversals $I = \{[1]_1, \dots, [1]_q\}$ and, for each $i \in \{1, \dots, q\}$, $T_i = (\{[-1]_1, \dots, [-1]_q\} - \{[-1]_i\}) \cup \{[1]_i\}$. As dependent transversals of a swirl-like matroid are circuit-hyperplanes, we can relax all the other circuit-hyperplanes of N_q to get W_q . As N_q was orientable, so is W_q by Theorem 6.1.1. \square

6.2.2.2. W_q is $GF(q)$ -representable and complex-representable.

Proof. Let α be a generator of $GF(q)$. Thus $\alpha^{(q-1)} = 1$ and $\alpha^j \neq 1$ for $j \in \{1, \dots, q - 2\}$. Consider the swirl-like matroid represented by the $GF(q)$ -matrix $A_q(\alpha)$. Note the labelling is an encoding with respect to this matrix. By choice of α , the only dependent transversal of this swirl-like matroid are $I = \{[1]_1, \dots, [1]_q\}$ and, for each $i \in \{1, \dots, q\}$, $T_i = (\{[\alpha]_1, \dots, [\alpha]_q\} - \{[\alpha]_i\}) \cup \{[1]_i\}$. Thus the swirl-like matroid that $A_q(\alpha)$ represents is W_q .

Similarly, if we take α to be a $(q - 1)$ -th primitive root of unity in \mathbb{C} , we have that $A_q(\alpha)$ is a complex-representation for W_q . \square

6.2.2.3. W_q is not real-representable.

Proof. Suppose W_q has a real-representation. By row operations and column scaling, there is an encoding with respect to the dependent transversal $I = \{[1]_1, \dots, [1]_q\}$, that is, each of $[1]_1, \dots, [1]_q$ is correctly labelled. We first see that the labelling $[x]_1, \dots, [x]_q$ is with respect to this encoding for some evaluation of x . The dependent transversals $T_i = (\{[x]_1, \dots, [x]_q\} - \{[x]_i\}) \cup \{[1]_i\}$ and $T_j = (\{[x]_1, \dots, [x]_q\} - \{[x]_j\}) \cup \{[1]_j\}$ only differ on the lines L_i and L_j . As all the terms in the one-products corresponding to T_i and T_j otherwise agree, this ensures that $[x]_j$ and $[x]_i$ encode the same value for every pair $i, j \in \{1, \dots, q\}$. Thus there is some real α that all $[x]_i$ encode, for the evaluation x to α . As I, T_1, \dots, T_q are the only dependent transversals, $\alpha^j \neq 1$ for $j \in \{1, \dots, q-2\}$ and $\alpha^{(q-1)} = 1$. However, as absolute value distributes over multiplication, the only real roots of unity are 1 and -1 , so this is a contradiction. \square

\square

Theorem 6.2.3. *For every positive integer k with $2^k - 1$ composite, there is a $GF(2^k)$ -representable, orientable matroid that is not real-representable.*

Proof. Say $2^k - 1 = st$ for odd integers $s, t > 1$ with s minimal. So s is prime. As in Section 3.4.2, let M_s be the rank- $(s+1)$ swirl-like matroid given by the $\mathbb{Q}(x)$ -matrix

$$A_s(x) = \begin{pmatrix} b_0 & [1]_0 & [x^{-1}]_0 & \dots & [x^{-s}]_0 & b_1 & [1]_1 & [x]_1 & [1]_2 & [x]_2 & \dots & [1]_s & [x]_s \\ \hline 1 & 1 & 1 & \dots & 1 & 0 & 0 & 0 & 0 & 0 & \dots & -1 & -x \\ 0 & -1 & -x^{-1} & \dots & -x^{-s} & 1 & 1 & 1 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & -1 & -x & 1 & 1 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & -1 & -x & \dots & 0 & 0 \\ \vdots & \vdots & \vdots & \ddots & \vdots & \vdots & \vdots & \vdots & \vdots & \vdots & \ddots & \vdots & \vdots \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 0 & 0 \\ 0 & 0 & 0 & \dots & 0 & 0 & 0 & 0 & 0 & 0 & \dots & 1 & 1 \end{pmatrix}$$

Note that the labellings are encoded with respect to the transversal $\{[1]_0, \dots, [1]_s\}$.

The dependent transversals of M_s are those whose labels multiply to 1, as in Section 2.2.2. This enforces representations that encode this labelling.

Let $z = x^s$, so we can also label the element $[x^{-s}]_0$ as $[z^{-1}]_0$. Now let R be the line restriction on the set $\{b_0, [1]_0, [z^{-1}]_0, b_1\}$. So that we may consider an amalgam across R , let M'_s be a copy of M_s where all the elements not in R have been renamed and where the indeterminant x has been replaced with y in labellings. Let L be the line with elements

$b_1 = [-\infty]_0$, $b_0 = [0]_0$, $[1]_0$, $[x^{-1}]_0, \dots, [x^{-(s-1)}]_0$, $[y^{-1}]_0, \dots, [y^{-(s-1)}]_0$, and $[z^{-1}]_0$ in that order. Note that for α with $0 < \alpha < 1$, the representation $A_s(\alpha)$ gives an orientation for M_s that is consistent with this ordering. Recalling that the indeterminant x has been replaced with y , the representation $A_s(\alpha)$ also gives a consistent orientation for M'_s . So by Remark 6.1.3, the proper amalgam $M_s \oplus_R M'_s$ is orientable.

6.2.3.1. $M_s \oplus_R M'_s$ is $GF(2^k)$ -representable and complex-representable.

Proof. Let α be a generator of $GF(2^k)^* \cong \mathbb{Z}_{st}$. Let $\beta = \alpha^{(t+1)}$ and $\gamma = \alpha^s = \beta^s$. Consider the homomorphism from $\mathbb{Q}(z, x, y)$ to $GF(2^k)$ by evaluating z, x, y as γ, α, β , respectively.

Let L_s and L'_s be the closures of L in M_s and M'_s respectively. As s is the minimal prime factor of $2^k - 1 = st$, if $i = j(t+1)$ in \mathbb{Z}_{st} then s divides $i - j$. So as α is a generator of $GF(2^k)^* \cong \mathbb{Z}_{st}$, this implies that L_s and L'_s have no coincident values besides those in L . Thus by Lemma 2.3.4, the proper amalgam $M_s \oplus_R M'_s$ is $GF(2^k)$ -representable.

Similarly, we can take α to be a $(q-1)$ -th primitive root of unity in \mathbb{C} , to get a complex-representation. \square

6.2.3.2. $M_s \oplus_R M'_s$ is not real-representable.

Proof. Suppose that $M_s \oplus_R M'_s$ has a real-representation. Say z is evaluated as γ on R . This representation induces encodings on the swirl-like matroid restrictions M_s and M'_s . Say x, y are evaluated as α, β in M_s and M'_s respectively. The labels encoded for elements of R in the restriction M_s enforces that α is the s^{th} root of γ . However, the restriction M'_s enforces that the s^{th} root of γ is β . However, as p is odd, γ only has one real s^{th} root by distribution of the absolute value over products. Thus $M_s \oplus_R M'_s$ is not real-representable. \square

Thus for a positive integer k where $2^k - 1$ is not prime and has minimal prime factor s , the matroid $M_s \oplus_R M_s$ is orientable and is representable over $GF(2^k)$ and the complex numbers, but not representable over the real numbers. \square

6.3 Which constructions still work?

As previously discussed, orientable matroids form a natural class [1, Proposition 7.9.1]. Thus, we can apply Theorems 4.1.1 and 5.6.1 to a complex-representable orientable matroid given by Theorem 6.2.1 to get the following theorems, respectively.

Theorem 6.3.1 (Strengthening of Theorem 4.0.2). *Each real-representable matroid is a minor of an excluded minor for real-representability that is complex-representable and orientable.*

Theorem 6.3.2 (Strengthening of Theorem 5.0.1). *Real-representability is not polynomially certifiable even within the class of complex-representability orientable matroids.*

The constructions used to prove Theorems 3.4.2 and 5.6.2 are non-real-representable matroids that are amalgamations of real-representable matroids. The real-representable constituents of these constructions are orientable by Theorem 6.0.3. However, to use Remark 6.1.3 to prove that the amalgam is orientable as well, we need to verify that each line across which amalgams occur have a consistent orientation. With the same evaluations of indeterminants in encodings to real transcendentals, we get an ordering of the common indeterminants and expressions of these indeterminants, and we can use a lexicographical ordering for other values. As described in Section 6.0.1, this will give a consistent orientation, as required. See Lemma 6.2.3 as an example. Once this is done, we have the following.

Theorem 6.3.3 (Strengthening of Theorem 3.4.2). *There is no sentence in the monadic second-order language MS_0 that characterizes real-representability for complex-representable orientable matroids.*

Theorem 6.3.4 (Strengthening of Theorem 5.6.2). *Non-real-representability is not polynomially certifiable even within the class of complex-representability orientable matroids.*

We have just made good use of how orientations of a line relate to orderings of values encoded on that line. It then seems presumptuous to imagine that we might be able to construct orientable matroids that are only representable over a given non-zero characteristic.

Conjecture 6.3.5 (Revision of Whittle's conjecture). *If an orientable matroid is representable over some field, then it is complex-representable.*

References

- [1] A. Björner, M. Las Vergnas, B. Sturmfels, N. White, G. Ziegler, *Oriented Matroids*, Second Edition, Cambridge University Press, Cambridge (1999).
- [2] R. Bland, M. Las Vergnas, Orientability of matroids, *J. Combinatorial Theory, Ser. B* **24** (1978), 94–123.
- [3] J. Bonin, Lattice path matroids: the excluded minors, *J. Combin. Theory Ser. B* **100**, (2010) 585–599.
- [4] G. Birkhoff, *Lattice Theory*, Third Edition, Amer. Math. Soc., Providence (1967).
- [5] T. Brylawski, An affine representation for transversal geometries, *Studies in Appl. Math.* **54**, (1975) 143–160.
- [6] T. Brylawski, Constructions. In *Theory of matroids* (ed. N. White), Cambridge University Press, Cambridge (1986), 127–223.
- [7] A. Cameron, D. Mayhew, Excluded minors for matroids satisfying Kinser’s inequalities, *Graphs Combin.* **32** (2016), 31–47.
- [8] R. Campbell, J. Geelen, On the complex-representable excluded minors for real-representability, *J. Combin. Theory Ser. B* (2019), 10.1016.
- [9] J. Edmonds, D. Fulkerson, Transversals and matroid partition, *J. Res. Nat. Bur. Standards Sect. B* **69B**, 147–153.
- [10] J. Edmonds, A. Mandel, Topology of oriented matroids, Ph.D. Thesis of A. Mandel, University of Waterloo.
- [11] J. Folkman, J. Lawrence, Oriented matroids, *J. Combin. Theory Ser. B* **25**, (1978) 199–236.
- [12] J. Geelen, Some open problems on excluding a uniform matroid, *Advances in Applied Mathematics* **41** (2008), 628–637.

- [13] J. Geelen, B. Gerards, G. Whittle, Solving Rota’s conjecture, *Notices Amer. Math. Soc.* **61** (2014), 736–743.
- [14] J. Geelen, G. Whittle, Certifying non-representability of matroids over prime fields, *J. Combin. Theory Ser. B* **117** (2016), 22–33.
- [15] P. 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), 470–479
- [16] W. Hochstättler, R. Nickel, Joins of oriented matroids, *Euro. J. Combin.* **32**, (2011) 841–852.
- [17] A. Ingleton, Representation of matroids, in: Combinatorial mathematics and its applications (Proc. Conf., Oxford, 1969), Academic Press, 1971, 149–167.
- [18] A. Ingleton, J. Piff, Gammoids and transversal matroids, *J. Combin. Theory Ser. B* **15** (1973), 51–68.
- [19] R. Kinser, New inequalities for subspace arrangements, *J. Combin. Theory Ser. A* **118** (2011), 152–161.
- [20] D. Knuth, The asymptotic number of geometries, *J. Combin. Theory Ser. A* **16** (1974), 398–400.
- [21] J. Kollár, Sharp effective nullstellensatz, *J. Amer. Math. Soc.* **4** (1988), 963–975.
- [22] T. Lazarsen, The representation problem for independence functions, *J. London Math. Soc.* **33** (1958), 21–25.
- [23] J. Lee, M. Scobee, A characterization of the orientations of ternary matroids, *J. Combin. Theory Ser. B* **77** (1999), 263–291.
- [24] S. MacLane, Some interpretations of abstract linear dependence in terms of projective geometry, *Amer. J. Math.* **58** (1936), 236–240.
- [25] J. Mason, On the class of matroids arising from paths in graphs, *Proc. London Math. Soc.* **25** (1972), 55–74.
- [26] F. Matúš, Classes of matroids closed under minors and principal extensions, *Combinatorica* **38** (2018), 935–954.
- [27] D. Mayhew, M. Newman, and G. Whittle, Yes, the “missing axiom” of matroid theory is lost forever, *Trans. Amer. Math. Soc.* **370** (2018), 5907–5929.
- [28] D. Mayhew, M. Newman, and G. Whittle, On excluded minors for real representability, *J. Combin. Theory Ser. B* **99** (2009) 685–689.

- [29] N. Mnëv, Varieties of combinatorial types of projective configurations and convex polyhedra, *Dokl. Akad. Nauk SSSR* **283** (1985), 1312–1314
- [30] E. Moore, A doubly-infinite system of simple groups, *Mathematical Papers Read at the International Mathematics Congress Held in Connection with the World's Columbian Exposition*, Macmillan & Co., 208–242.
- [31] T. Nakasawa. Zur Axiomatik der linearen Abhängigkeit. I, *Sci. Rep. Tokyo Bunrika Daigaku Sect. A* **2** (1935), 129–149.
- [32] P. Nelson, Almost all matroids are nonrepresentable, *Proc. London Math. Soc.* **50** (2018), 245–248.
- [33] A. Nerode, Linear automaton transformations, *Proc. Amer. Math. Soc.* **9** (1958), 541–544.
- [34] J. Oxley, Matroid Theory, Second Edition, Oxford University Press, New York (2011).
- [35] J. Oxley, A matroid extension result, *SIAM J. Discrete Math.* **33** (2019), 138–152.
- [36] P. D. Seymour, Recognizing graphic matroids, *Combinatorica* **1** (1981), 75–78.
- [37] P. D. Seymour, Decomposition of regular matroids, *J. Combin. Theory Ser. B* **28** (1980), 305–359.
- [38] K. von Staudt, *Beitrage zur Geometrie der Lage Vol. 2*, (1957) Nürnberg.
- [39] B. Sturmfels, On the decidability of Diophantine problems in combinatorial geometry, *Bull. Amer. Math. Soc.* **35** (1987), 314–326.
- [40] A. Tarski, A decision method for elementary algebra and geometry, Manuscript. Santa Monica, CA: RAND Corp., (1948).
- [41] W. Tutte, A homotopy theorem for matroids. I, II, *Trans. Amer. Math. Soc.* **88** (1958), 144–174.
- [42] W. Tutte, Lectures on matroids, *J. Res. Nat. Bur. Standards Sect. B* **69B** (1965), 1–47.
- [43] P. Vámos, The missing axiom of matroid theory is lost forever, *J. London Math. Soc.* **18** (1978), 403–408.
- [44] N. White, ed. Combinatorial Geometries, Cambridge University Press, (1987).
- [45] H. Whitney, On the abstract properties of linear dependence, *Amer. J. Math.* **57** (1935), 509–533.