Solving Rota's Conjecture

Jim Geelen, Bert Gerards, and Geoff Whittle

In 1970, Gian-Carlo Rota posed a conjecture predicting a beautiful combinatorial characterization of linear dependence in vector spaces over any given finite field. We have recently completed a fifteen-year research program that culminated in a solution of Rota's Conjecture. In this article we discuss the conjecture and give an overview of the proof.

Matroids are a combinatorial abstraction of linear independence among vectors; given a finite collection of vectors in a vector space, each subset is either dependent or independent. A matroid consists of a finite *ground set* together with a collection of subsets that we call *independent*; the independent sets satisfy natural combinatorial axioms coming from linear algebra. Not all matroids can be represented by a collection of vectors and, ever since their introduction by Hassler Whitney [26] in 1935, mathematicians have sought ways to characterize those matroids that are. Rota's Conjecture asserts that representability over any given *finite field* is characterized by a finite list of obstructions. We will formalize these notions, and the conjecture, in the next section. In the remainder of this introduction, we will describe the journey that led us to a solution.

In the late 1990s, Rota's Conjecture was already known to hold for fields of size two, three, and four. Geelen, Gerards, and Kapoor [4] had recently announced that there are seven obstructions for representability over the four element field. As far back as 1958, Tutte had already proved that there is one obstruction for the class of binary matroids [22], and in the 1970s Bixby [1] and Seymour [20] had independently proved that there are four obstructions for the class of ternary matroids, verifying a result that was announced earlier by Ralph Reid.

It was a promising time for matroid theory; a number of useful new techniques had emerged in the previous decade, and it looked like there was real potential for major progress. Nevertheless, there is a significant difference between the concrete problem of finding the full set of obstructions for some particular field and the abstract problem of showing that there are finitely many obstructions for an arbitrary finite field. The techniques available at that time did not offer any realistic hope of solving Rota's Conjecture in general.

Our hopes were raised in 1999, at a workshop on graph theory in Oberwolfach, Germany, when Neil Robertson and Paul Seymour proposed ideas for extending their Graph Minors Project to matroids. We eagerly took up the challenge; the Graph Minors Theory had exactly the kind of general purpose tools that we lacked. Their results were published in a series of twenty-three journal papers totalling more than 700 pages, and we spent the next year learning how the machinery worked.

Then we set about the daunting task of extending the Graph Minors Theory to matroids. That was a major undertaking which took a decade to complete. The first three years were the hardest, and during that time it was not at all clear that we would succeed. The relief was palpable when we made our first major breakthrough in November of 2002. That marked a massive turning point in the project: from that point on we had no doubt that we would get to the end, despite the enormity

Jim Geelen is a professor at the University of Waterloo, Canada. His email address is jim.geelen@uwaterloo.ca.

Bert Gerards is a scientific staff member at the Centrum Wiskunde & Informatica in Amsterdam, The Netherlands. His email address is bert@gerardsbase.nl.

Geoff Whittle is a professor at Victoria University of Wellington, New Zealand. His email address is geoff.whittle@vuw.ac.nz.

of the task that lay ahead. We even maintained our optimism through two bleak years, 2006 and 2007, when we made almost no discernible progress despite considerable effort.

In early 2011 we had extended the main part of the Graph Minors Project to matroids. The techniques that we had developed were general and powerful, however, there were still many issues that were particular to Rota's Conjecture that had to be addressed. In fact, at that time we did not have a concrete approach in mind, and our first serious attempt to prove Rota's Conjecture, later in 2011, failed. In early 2012 we hatched a new plan of attack and spent the rest of that year developing the required machinery. When we met together in January 2013, we were confident that we had the right method and all of the required tools; after three weeks of joint work we had thrashed out the remaining details.

We are now immersed in the lengthy task of writing up our results. Since that process will take a few years, we have written this article offering a high-level preview of the proof.

What Is a Matroid?

A *matroid* consists of a pair $(E, \mathcal{I})$ where E is a finite set, called the *ground set*, and $\mathcal{I}$ is a collection of subsets of E, called *independent sets*, such that

1. the empty set is independent,
2. subsets of independent sets are independent, and
3. for each set $X \subseteq E$, the maximal independent subsets of X all have the same size.

The canonical example of a matroid comes from a collection of vectors in a vector space or, equivalently, the columns of a matrix. Let A be a matrix over a field $\mathbb{F}$, and let E be the set of column-indices of A. The *column matroid* of A, denoted by $M(A)$, is the pair $(E, \mathcal{I})$ where $\mathcal{I}$ is the collection of subsets of E that index linearly independent sets of columns. A matroid is called $\mathbb{F}$-*representable* if it is the column matroid of a matrix over the field $\mathbb{F}$.

The matroid axioms allow us to extend notions such as *bases* and *rank* from linear algebra to matroids; a *basis* being a maximal independent set and the *rank* of a set X in a matroid $M = (E, \mathcal{I})$, denoted $r_M(X)$, being the size of a maximal independent subset of X.

There are matroids that are not representable over any field; indeed, it is believed, but not yet proven, that the proportion of n-element matroids that are representable is vanishingly small as n tends to infinity. It was Whitney himself, in the introductory paper on matroid theory [26] in 1935, who posed the problem of characterizing the class of representable matroids.

Whitney's problem is open to interpretation in that there are many different types of characterizations that one might consider and that one can consider representability over a specified or unspecified field. With two notable exceptions, most interpretations of Whitney's problem have been met with negative answers; see [11], [12], [13], [21]. The first exception is the algorithmic problem of determining when a given matroid is representable over an unspecified field, which was proved to be decidable by Vámos [28]. The second exception is, of course, Rota's Conjecture.

Matroid Duality

Matroids come in dual pairs; the *dual* of M, denoted M^*, is the matroid on E whose bases are the complements of the bases of M. While it is not immediately clear from this definition, M^* is indeed always a matroid. More surprising, perhaps, is the fact that $\mathbb{F}$-representability is preserved under duality. Indeed, if $M = M(A)$, then $M(A)^* = M(A')$, where A' is a matrix whose row space is the orthogonal space of the row space of A.

Minors of Matroids

Let C and D be sets of elements in a matroid $M = (E, \mathcal{I})$. The matroid obtained from M by *deleting* D is defined as $(E - D, \{I \subseteq E - D : I \in \mathcal{I}\})$. Note that $M \setminus D$ is clearly always a matroid. The dual operation of deletion is *contraction* which is defined by $M/C = (M^* \setminus C)^*$. This definition of contraction is not particularly illuminating; geometrically it corresponds to *projection* from C, which is more readily seen via the rank function

$$r_{M/C}(X) = r_M(X \cup C) - r_M(C).$$

A *minor* of M is a matroid of the form $M \setminus D/C$, where D and C are disjoint subsets of E; when $D \cup C$ is nonempty, we call the minor *proper*. Note that the class of $\mathbb{F}$-representable matroids is closed under both deletion and duality and, hence, also under taking minors. For any minor-closed class of matroids, it is natural to consider characterizing the class by describing the *excluded minors*, that is, the matroids outside the class whose proper minors are all in the class—which brings us to Rota's Conjecture [19].

Conjecture 1 (Rota). *For each finite field $\mathbb{F}$, there are, up to isomorphism, only finitely many excluded minors for the class of $\mathbb{F}$-representable matroids.*

Since a minor-closed class is determined by its list of excluded minors, Rota's Conjecture provides a succinct characterization for the class of $\mathbb{F}$-representable matroids.

Figure 1. The 4-point line, $U_{2,4}$; the 5-point line, $U_{2,5}$; and the Fano plane, F_7.

Figure 2. The Kuratowski graphs: K_5 and $K_{3,3}$.

Combinatorial Geometry

Consider a representable matroid $M = M(A)$. If A' is a matrix that is obtained from A by elementary row operations and column scaling, then A' is also a representation. Observe that, in obtaining A' from A in this way, we have simply applied a projective transformation to the column space of A, and, hence, we say that A and A' are *projectively equivalent*.

Now, instead of considering the elements in a representable matroid as vectors in a vector space, we could equally well regard them as points in a projective geometry. While this is a relatively small shift in perspective, it is surprisingly useful for intuition, and a lot of matroid terminology comes from geometry. For example, a *rank-k flat* is a maximal set of rank k. A *point* is a rank-1 flat, a *line* is a rank-2 flat, a *plane* is a rank-3 flat, and a *hyperplane* is a flat of rank $r_M(E) - 1$.

This geometric approach has the additional feature that it provides a succinct method for describing matroids of low rank. For example, Figure 1 depicts three important matroids $U_{2,4}$, $U_{2,5}$, and F_7. The 4-point line, $U_{2,4}$, is the only excluded minor for the class of binary matroids, and the excluded minors for the class of ternary matroids are $U_{2,5}$, F_7, and their duals.

Graph Minors Project

A *minor* of a graph G is a graph that is obtained from a subgraph of G by *contracting* some edges. To contract an edge, one identifies the ends of the edge into a single vertex and then deletes the edge.

Note that the class of planar graphs is closed under taking minors. As we will see, this class turns out to be particularly fundamental in the study of graph minors. It is therefore fitting to begin with Kuratowski's characterization of the class of planar graphs [10], which is one of the gems of mathematics.

Theorem 2 (Kuratowski's Theorem). *A graph is not planar if and only if it has a minor isomorphic to $K_{3,3}$ or K_5.*

We can restate Kuratowski's Theorem as: the excluded minors for the class of planar graphs are $K_{3,3}$ and K_5; see Figure 2.

Robertson and Seymour generalized this result from the plane to arbitrary surfaces in [16].
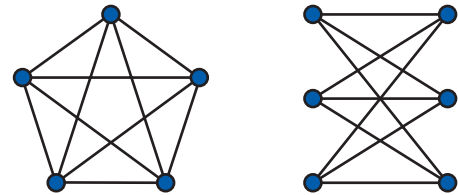
Theorem 3 (Generalized Kuratowski's Theorem). *For any given surface, there are only finitely many excluded minors for the class of graphs that embed in the surface.*

Later Robertson and Seymour further generalized the result to arbitrary minor-closed classes [18]. Diestel, in his book on graph theory [2], says that this theorem "dwarfs any other result in graph theory and may doubtless be counted among the deepest theorems that mathematics has to offer".

Theorem 4 (Well-Quasi-Ordering Theorem). *Each minor-closed class of graphs has only finitely many excluded minors.*

The following two statements are reformulations of the WQO Theorem:

- In each infinite set of graphs there are two graphs, one isomorphic to a minor of the other.
- There are only countably many distinct minor-closed classes of graphs.

How then does one go about proving such a result? The first step is obvious. Suppose that the result is false and, hence, there exists an infinite sequence $(H_1, H_2, \ldots)$ of graphs such that none is isomorphic to a minor of another. In particular, none of the graphs $(H_2, H_3, \ldots)$ contains a minor isomorphic to H_1. The next step is to ask what a graph looks like if it does not contain a minor isomorphic to H_1. This question is answered by the Graph Minors Structure Theorem [17], which is the workhorse of the Graph Minors Project.

Let $EX(H_1)$ denote the set of graphs that do not contain a minor isomorphic to H_1. If S is a surface into which H_1 does not embed, then every graph that embeds in S is clearly in $EX(H_1)$. Let $\mathcal{G}$ denote the set of graphs that embed into a surface that does not embed H_1. Thus $\mathcal{G} \subseteq EX(H_1)$. The gist of the Graph Minors Structure Theorem is that each graph in $EX(H_1)$ can be constructed from graphs in $\mathcal{G}$ in a specified way.

Note that, if H_1 is planar, then it embeds in every surface and, hence, the class $\mathcal{G}$ is empty. Therefore graphs omitting H_1 have a particularly simple structure; see [15]. The extension of that result to matroids over finite fields [5] was our first major breakthrough.

The Graph Minors Structure Theorem gives considerable traction with the WQO Theorem, but completing the proof still requires more than 100 pages of careful argument.

From Graphs to Matroids

A *circuit* in a graph $G = (V, E)$ is a connected subgraph of G whose vertices all have degree 2, and a *forest* of G is a subgraph of G that does not contain a circuit. The *cycle matroid* of G, denoted $M(G)$, is defined as $(E, \mathcal{F})$ where $\mathcal{F}$ is the collection of all edge-sets of forests of G. It is easy to show that $M(G)$ is always a matroid. A matroid is called *graphic* if it is the cycle matroid of a graph.

The class of graphic matroids turns out to be surprisingly fundamental, in much the same way that the class of planar graphs is fundamental in the Graph Minors Project. Substantial parts of matroid theory have their roots in graph theory, which, given what we now know, is not at all surprising. This link to graph theory has also shaped the terminology. Note that the minimal dependent sets of $M(G)$ are the edge-sets of circuits of G. This terminology is carried over to general matroids: a minimal dependent set in a matroid is called a *circuit*.

Minors and Duality

The class of graphic matroids is minor-closed. Moreover, if D and C are disjoint edge-sets in a graph G, then $M(G \setminus D/C) = M(G) \setminus D/C$. The class of graphic matroids is not, however, closed under duality. In fact, a graph G is planar if and only if $M(G)^*$ is graphic. This beautiful geometric characterization of planarity was observed by Whitney [25] and is the first glimpse of a deep connection between matroid theory and topological graph theory.

By the discussion above, any characterization of the class of graphic matroids gives a characterization for the class of planar graphs. In particular, the following excluded-minor characterization for the class of graphic matroids, due to Tutte [23], implies Kuratowski's Theorem.

Theorem 5. *The excluded minors for the class of graphic matroids are $U_{2,4}$, F_7, $(F_7)^*$, $M(K_{3,3})^*$, and $M(K_5)^*$.*

Representability

Graphic matroids are representable; in fact, they are representable over every field. The *signed-incidence* matrix of a graph $G = (V, E)$ is a matrix $A \in \{0, \pm 1\}^{V \times E}$ such that, for each edge e with ends u and v, if $u = v$, then the column indexed by e is identically zero, and if $u \neq v$, then the column indexed by e has exactly one 1 and exactly one -1, and these are in the (u, e) and (v, e) positions, in
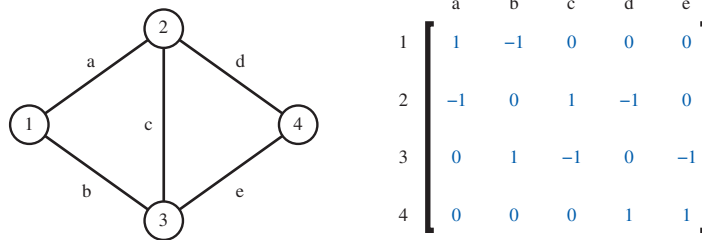

Figure 3. A graph and its signed-incidence matrix.

no particular order; see Figure 3. It is a routine exercise to verify that A is a representation of $M(G)$ over any field.

Frame Matroids

The class of graphic matroids is contained in a richer class of matroids, namely the class of "frame matroids". A *frame matrix* over $\mathbb{F}$ is a matrix with at most two nonzero entries per column, and a *frame matroid* over $\mathbb{F}$ is one that is represented by a frame matrix. As with the class of graphic matroids, the class of frame matroids is closed under taking minors. Moreover, there is a natural way to associate a graph with a frame matroid, although, to fully describe the matroid one needs to "decorate" the graph by orienting and labelling its edges.

Matroid Minors Project

Rota's Conjecture is reminiscent of the Generalized Kuratowski Theorem. Robertson and Seymour were able to further generalize the Generalized Kuratowski Theorem to obtain the WQO Theorem, so, perhaps, Rota's Conjecture is a special case of a much more general theorem. However, it is not true that the WQO Theorem extends to all matroids. An *antichain* of matroids is a set of matroids no one of which is isomorphic to a minor of another. It is surprisingly easy to construct infinite antichains; the matroids $\{M_3, M_4, M_5, \ldots\}$ in Figure 4, for example, form an infinite antichain of rank-3 matroids.

We were, however, able to prove the following significant generalization of the WQO Theorem.

Theorem 6 (Matroid WQO Theorem). *For each finite field $\mathbb{F}$ and each minor-closed class of $\mathbb{F}$-representable matroids, there are only finitely many $\mathbb{F}$-representable excluded minors.*

This result is incomparable with Rota's Conjecture, in the sense that neither result implies the other, since, in Rota's Conjecture, the excluded minors are, by definition, not representable over $\mathbb{F}$.

Our proof of the Matroid WQO Theorem parallels the proof of the WQO Theorem for graphs. The
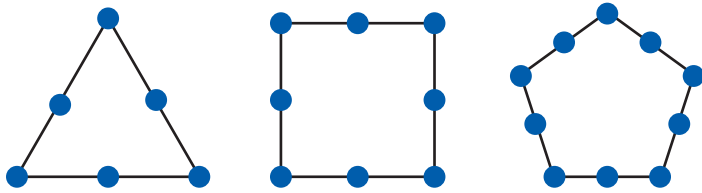
Figure 4. The first three matroids, M_3, M_4, and M_5, in an infinite antichain.

most significant ingredient is an analogue of the Graph Minors Structure Theorem. Recall that the Graph Minors Structure Theorem essentially says that, for any minor closed class of graphs, the graphs in the class can be constructed, in specified ways, from graphs that embed in surfaces of low genus.

For a finite field $\mathbb{F}$, what are the fundamental minor-closed classes of $\mathbb{F}$-representable matroids? It turns out that there are, up to duality, only two types. First, for each subfield $\mathbb{F}'$ of $\mathbb{F}$, the set of all $\mathbb{F}'$-representable matroids is minor-closed. The second is the class of frame matroids over $\mathbb{F}$.

The Matroid Minors Structure Theorem describes, for a finite field $\mathbb{F}$, how to construct, in a specified way, the members of an arbitrary fixed minor-closed class of $\mathbb{F}$-representable matroids from matroids represented over subfields, frame matroids, and the duals of frame matroids. It is this theorem that consumed most of our effort—it took more than a decade to prove.

Toward Rota's Conjecture

While we have been working together since 1999, it was not until 2012 that we finally formulated our strategy for attacking Rota's Conjecture. Prior to 2012 we were primarily immersed in the Matroid Minors Structure Theory. In this section we outline the machinery that we developed specifically for proving Rota's Conjecture. Significant parts of the machinery were developed with Tony Huynh and Stefan van Zwam who were postdoctoral fellows with Bert Gerards.

Connectivity

Connectivity plays a significant role in matroid theory in general and in the proof of Rota's Conjecture in particular. The following definitions are due to Tutte [24]. A *k-separation* in a matroid $M = (E, \mathcal{I})$ is a partition (X, Y) of E such that $r_M(X) + r_M(Y) - r_M(E) < k$ and $|X|, |Y| \geq k$. A matroid is *k-connected* if it has no *l*-separation for any $l < k$.

To motivate the definition, consider the case that M is represented by a matrix A. Let $\mathbb{V}$ denote the column space of A, and let $\langle X \rangle$ denote the subspace of $\mathbb{V}$ spanned by the columns indexed by X. The dimension theorem for subspaces shows that the dimension of $\langle X \rangle \cap \langle Y \rangle$ is equal to $r_M(X) + r_M(Y) - r_M(E)$. So when $r_M(X) + r_M(Y) - r_M(E)$ is small, X and Y are separated by a low-dimensional subspace of $\mathbb{V}$. The condition that $|X|, |Y| \geq k$ is a nondegeneracy condition that is needed since $r_M(X) + r_M(Y) - r_M(E) \leq \min(|X|, |Y|)$.

Problems concerning matroid representation typically reduce to the 3-connected case, and this is true of Rota's Conjecture.

Lemma 7. *For each field $\mathbb{F}$, each excluded minor for the class of $\mathbb{F}$-representable matroids is 3-connected.*

Unfortunately one cannot easily improve on Lemma 7; for each finite field of order at least 3, there is an excluded minor that is not 4-connected. The problem is that the nondegeneracy condition $|X|, |Y| \geq k$ is too weak. We were able to prove:

Theorem 8. *For each finite field $\mathbb{F}$ and integer k there is an integer $n = n(\mathbb{F}, k)$ such that, if (A, B) is a k-separation in an excluded minor for the class of $\mathbb{F}$-representable matroids, then $\min(|A|, |B|) \leq n$.*

We do not know an elementary proof of Theorem 8; our proof relies on several difficult results, including the Matroid WQO Theorem. As a consequence of applying well-quasi-ordering techniques in the proof, we do not have a computable bound on $n(\mathbb{F}, k)$, which, in turn, means that we do not obtain a computable bound on either the number or the size of the excluded minors for the class of $\mathbb{F}$-representable matroids.

Theorem 8 shows that the excluded minors are very highly connected in some sense. Henceforth we will focus on excluded minors that are very highly connected in Tutte's sense. While we lose some generality in doing so, we avoid a lot of technicalities without bypassing the main issues.

Inequivalent representations

Recall that, if A_1 and A_2 are projectively equivalent matrices, then they represent the same matroid. There are matroids that admit many projectively inequivalent representations over a given finite field; this is one of the major sources of difficulty in the area of matroid representation and is the focus of much attention [6], [8], [9], [14], [27]. This difficulty is tamed by considering matroids that have sufficiently high connectivity.

Theorem 9. *For each finite field $\mathbb{F}$, there exist integers k_1 and n_1 such that each k_1-connected matroid has at most n_1 projectively inequivalent representations over $\mathbb{F}$.*

Theorem 9 was proved in 2012 by Geelen, Gerards, Huynh, and Van Zwam [3]. That paper develops groundbreaking inductive tools for dealing with highly-connected matroids. The key is to

weaken the notion of connectivity in a manner compatible with Theorem 8. Those connectivity tools should find wider application in combinatorics, particularly in graph theory.

Bifurcation of Representations and Fixed Elements

Let $\mathbb{F}$ be a finite field, and let k_1 and n_1 be the integers given by Theorem 9. Suppose that e is an element of an $\mathbb{F}$-representable matroid and that M and $M \setminus e$ are both k_1-connected. While Theorem 9 shows that M and $M \setminus e$ both have at most n_1 inequivalent representations, it does not describe how the representations are related. Some representations of $M \setminus e$ may not extend to representations of M while others may bifurcate to two or more inequivalent representations of M; this bifurcation is a significant cause of difficulty for us. Fortunately we understand the root cause of bifurcation.

Suppose that A is a representation of $M \setminus e$ that extends to two projectively inequivalent representations $[A, u]$ and $[A, v]$ of M. Let

$$A' = \begin{pmatrix} e & e' \\ A & u & v \end{pmatrix},$$

and let $M' = M(A')$. Note that swapping e and e' defines an automorphism of M'; we say that e and e' are *clones* in M'. Moreover, since $[A, u]$ and $[A, v]$ are projectively inequivalent, $\{e, e'\}$ is an independent set in M'. We say that an element e is *fixed* in M if it is not possible to get a matroid M' by extending M by an element e' so that $\{e, e'\}$ is an independent clonal pair. The next lemma follows directly from the discussion above.

Lemma 10. *If e is a fixed element in a matroid M, then no representation of $M \setminus e$ extends to two projectively inequivalent representations of M.*

The following result shows that highly connected $\mathbb{F}$-representable matroids contain few elements that are not fixed.

Theorem 11. *For each finite field $\mathbb{F}$, there exist integers k_2 and n_2 such that each k_2-connected $\mathbb{F}$-representable matroid has at most n_2 elements that are not fixed.*

Theorem 11 was proved by Geelen and Van Zwam [7]. The proof relies on the main result of [5], which describes the structure of an $\mathbb{F}$-representable matroid when we exclude the cycle matroid of a planar graph as a minor.

Theorems 9 and 11 provide us with a lot of control over the inequivalent representations of a matroid. However, as early as 1999, we already had much stronger results controlling the inequivalent representations of GF(5)-representable matroids [14], [27], and yet we still do not know the full set of excluded minors.
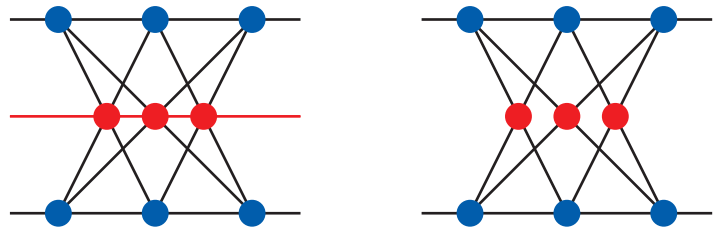

Figure 5. The Pappus matroid and the non-Pappus matroid.

Relaxing a Circuit-Hyperplane

Recall that a hyperplane in a matroid $M = (E, \mathcal{I})$ is a maximal set of rank $r_M(E) - 1$ and a circuit is a minimal dependent set. If C is both a circuit and a hyperplane, then we can construct a new matroid M' from M by *relaxing* C, that is $M' = (E, \mathcal{I} \cup \{C\})$.

This construction does not respect representability. Consider for example the matroids in Figure 5. The Pappus matroid is representable over the reals, which is evident by its straight-line drawing. The *non-Pappus* matroid, is obtained from the Pappus matroid by relaxing the circuit-hyperplane indicated in red. Pappus's *Collection* (circa 340 A.D.) already contained a proof that the non-Pappus matroid is not representable over the reals. In fact, the non-Pappus matroid is not representable over any field.

While the operation of relaxing a circuit-hyperplane does not behave well with respect to representation in general, it behaves particularly poorly with respect to representation over finite fields. By making this explicit, we obtain a method for certifying nonrepresentability in the proof of Rota's Conjecture. The following result shows that, if M is a sufficiently large and sufficiently connected matroid, and if M' is a matroid obtained from M by relaxing a circuit-hyperplane, then M and M' cannot both be representable over a given finite field.

Theorem 12. *For each finite field $\mathbb{F}$, there exist integers k_3 and n_3 such that, if M_1 and M_2 are $\mathbb{F}$-representable matroids on a common ground set E, where M_2 is obtained from M_1 by relaxing a circuit-hyperplane and M_1 is k_3-connected, then $|E| \leq n_3$.*

In fact, we proved a stronger result. Note that, if M_2 is obtained from M_1 by relaxing C, then $M_1 \setminus e = M_2 \setminus e$ for each $e \in C$.

Theorem 13. *For each finite field $\mathbb{F}$, there exist integers k_4 and n_4 such that, if M_1 and M_2 are distinct $\mathbb{F}$-representable matroids on a common ground set E and M_1 is k_4-connected, then there is a set $X \subseteq E$ such that $|X| \leq n_4$ and $M_1 \setminus e \neq M_2 \setminus e$ for each $e \in E - X$.*

This result is the hardest step in resolving Rota's Conjecture and relies on the full strength of the Matroid Minors Structure Theory that we had developed earlier. Much of the proof is graph theoretic, which is a little surprising given the statement of the result.

Highly Connected Excluded Minors

In this final section, we indicate how the machinery developed in the previous section can be combined to prove a weakening of Rota's Conjecture. In particular, we prove that, for each finite field $\mathbb{F}$, there are only finitely many highly connected excluded minors for the class of $\mathbb{F}$-representable matroids. This proof exemplifies all of the techniques needed to prove Rota's Conjecture in general other than one technical connectivity result which we discuss after the proof.

Theorem 14. *For each finite field $\mathbb{F}$, there exists an integer k such that the class of $\mathbb{F}$-representable matroids has no k-connected excluded minor with $2k$ or more elements.*

Proof. Let k_1, k_2, k_4, n_1, n_2, and n_4 be the integers given by Theorems 9, 11, and 13. Let

$$t = n_1(n_2 + 1)(n_4 + 1),$$

and let

$$k = \max(k_1 + n_1(n_4 + 1), k_2 + t, k_4 + 1).$$

Suppose that there is a k-connected excluded minor $M = (E, \mathcal{I})$ with $|E| > 2k$.

Throughout the proof we freely make use of the fact that, for any set $Z \subseteq E$, $M \setminus Z$ is $(k - |Z|)$-connected.

Let X_1 be a t-element subset of E, and let Π be a partition of X into $n_1(n_4 + 1)$ sets of size $n_2 + 1$. By Theorem 11, for each $Y \in \Pi$, there is an element $e \in Y$ that is fixed in $M \setminus (X_1 - Y)$. Let X_2 be a transversal of Π containing one such element from each set in Π; thus $|X_2| = n_1(n_4 + 1)$.

By Theorem 9, $M \setminus X_2$ has at most n_1 inequivalent representations. Since M is an excluded minor, $M \setminus e$ is representable for each $e \in X_2$. So there exists a representation A_0 of $M \setminus X_2$ and an $(n_4 + 1)$-element subset X_3 of X_2 such that A_0 extends to a representation of $M \setminus e$ for each $e \in X_3$. By construction, each element $e \in X_2$ is fixed in $M \setminus (X_2 - \{e\})$. So, by Lemma 10, the representation A_0 extends uniquely to a representation of $M \setminus (X_2 - \{e\})$. Thus we obtain a unique matrix A such that $A|(E - X_2) = A_0$ and, for each $e \in X_3$, $M(A) \setminus e = M \setminus e$.

Since M is not $\mathbb{F}$-representable, $M \neq M(A)$. Therefore there exists a set $B \subseteq E$ that is a basis in exactly one of M and $M(A)$. Note that $X_3 \subseteq B$. Choose $f \in E - B$. Now $M \setminus f$ and $M(A) \setminus f$ are distinct $\mathbb{F}$-representable matroids on the same ground set. Moreover $M \setminus f, e = M(A) \setminus f, e$ for

each $e \in X_3$. This contradicts Theorem 13, and this contradiction completes the proof. $\qquad\square$

To prove Rota's Conjecture in general, we start with an n-element excluded minor M, as in the proof above. By Theorem 8, M is "weakly" k-connected. Then, using the techniques developed in [3], we find a t-element set X_1 such that, after possibly replacing M with its dual, $M \setminus D$ remains weakly k-connected for each $D \subseteq X_1$. The rest of the proof continues as above.

References

[1] R. E. BIXBY, On Reid's characterization of ternary matroids, *J. Combin. Theory Ser. B* **26** (1979), 174–204.
[2] R. DIESTEL, *Graph theory*, Springer-Verlag, New York, 1997.
[3] J. GEELEN, B. GERARDS, T. HUYNH, and S. VAN ZWAM, Generating k-connected matroids (working title), in preparation.
[4] J. F. GEELEN, A. M. H. GERARDS, and A. KAPOOR, The excluded minors for GF(4)-representable matroids, *J. Combin. Theory Ser. B* **79** (2000), 247–299.
[5] J. GEELEN, B. GERARDS, and G. WHITTLE, Excluding a planar graph from GF(q)-representable matroids, *J. Combin. Theory Ser. B* **97** (2007), 971–998.
[6] ———, On inequivalent representations of matroids over nonprime fields, *J. Combin. Theory Ser. B* **100** (2010), 740–743.
[7] J. GEELEN and S. VAN ZWAM, Fixed elements and matroid tangles (working title), in preparation.
[8] J. GEELEN and G. WHITTLE, Inequivalent representations of matroids over prime fields, *Adv. in Appl. Math.* **51** (2013), 1–175.
[9] J. KAHN, On the uniqueness of matroid representations over GF(4), *Bull. London Math. Soc.* **20** (1988), 5–10.
[10] K. KURATOWSKI, Sur le problème des courbes gauches en topologie, *Fund. Math.* **15** (1930), 271–283.
[11] T. LAZARSON, The representation problem for independence functions, *J. London Math. Soc.* **33** (1958), 21–25.
[12] D. MAYHEW, M. NEWMAN, and G. WHITTLE, On excluded minors for real-representable matroids, *J. Combin. Theory Ser. B* **99** (2009), 685–689.
[13] ———, Is the missing axiom of matroid theory lost forever?, arXiv:1204.3365v1.
[14] J. OXLEY, D. VERTIGAN, and G. WHITTLE, On inequivalent representations of matroids over finite fields, *J. Combin. Theory Ser. B* **67** (1996), 325–343.
[15] N. ROBERTSON and P. D. SEYMOUR, Graph Minors. IV. Excluding a graph, *J. Combin. Theory Ser. B* **48** (1990), 227–254.
[16] ———, Graph Minors. VIII. A Kuratowski theorem for general surfaces, *J. Combin. Theory Ser. B* **48** (1990), 255–288.
[17] ———, Graph minors. XVI. Excluding a nonplanar graph, *J. Combin. Theory Ser. B* **89** (2003), 43–76.
[18] ———, Graph Minors. XX. Wagner's Conjecture, *J. Combin. Theory Ser. B* **92** (2004), 325–357.
[19] G.-C. ROTA, *Combinatorial theory, old and new*, In Proc. Internat. Cong. Math. (Nice, 1970), pp. 229–233. Gauthier-Villars, Paris.

[20] P. D. Seymour, Matroid representation over GF(3), *J. Combin. Theory Ser. B* **26** (1979), 159–173.

[21] ———, Recognizing graphic matroids, *Combinatorica* **1** (1981), 387–394.

[22] W. T. Tutte, A homotopy theorem for matroids, I, II. *Trans. Amer. Math. Soc.* **88** (1958), 144–174.

[23] ———, Matroids and graphs, *Trans. Amer. Math. Soc.* **90** (1959), 527–552.

[24] ———, Lectures on matroids, *J. Nat. Bur. Standards Sect. B* **69B** (1965), 1–47.

[25] H. Whitney, Nonseparable and planar graphs, *Trans. Amer. Math. Soc.* **34** (1932), 339–362.

[26] ———, On the abstract properties of linear independence, *Amer. J. Math.* **57** (1935), 509–533.

[27] G. Whittle, Stabilizers of classes of representable matroids, *J. Combin. Theory Ser. B* **77** (1999), 39–72.

[28] P. Vámos, *A necessary and sufficient condition for a matroid to be linear*, In Möbius algebras (Proc. Conf. Univ. Waterloo, 1971), 162–169. University of Waterloo, Waterloo.
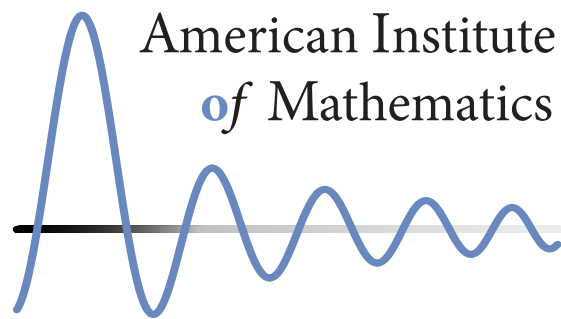
American Institute of Mathematics

AIM, the American Institute of Mathematics, sponsors week-long activities in all areas of the mathematical sciences with an emphasis on focused collaborative research.

Call for Proposals

Workshop Program

AIM invites proposals for its focused workshop program. AIM's workshops are distinguished by their specific mathematical goals. This may involve making progress on a significant unsolved problem or examining the convergence of two distinct areas of mathematics. Workshops are small in size, up to 28 people, to allow for close collaboration among the participants.

SQuaREs Program

AIM also invites proposals for the SQuaREs program: Structured Quartet Research Ensembles. More long-term in nature, this program brings together groups of four to six researchers for a week of focused work on a specific research problem in consecutive years.

More details are available at:

http://www.aimath.org/research/

deadline: November 1

AIM seeks to promote diversity in the mathematics research community. We encourage proposals which include significant participation of women, underrepresented minorities, junior scientists, and researchers from primarily undergraduate institutions.