Matroids with nine elements

Dillon Mayhew
School of Mathematics, Statistics & Computer Science
Victoria University
Wellington, New Zealand
Dillon.Mayhew@mcs.vuw.ac.nz

Gordon F. Royle
School of Computer Science & Software Engineering
University of Western Australia
Perth, Australia
gordon@csse.uwa.edu.au

February 13, 2007

Abstract

We describe the computation of a catalogue containing all matroids with up to nine elements, and present some fundamental data arising from this cataogue. Our computation confirms and extends the results obtained in the 1960s by Blackburn, Crapo & Higgs. The matroids and associated data are stored in an online database, and we give three short examples of the use of this database.

1 Introduction

In the late 1960s, Blackburn, Crapo & Higgs published a technical report describing the results of a computer search for all simple matroids on up to eight elements (although the resulting paper [2] did not appear until 1973). In both the report and the paper they said

"It is unlikely that a complete tabulation of 9-point geometries will be either feasible or desirable, as there will be many thousands of them. The recursion $g(9) = g(8)^{3/2}$ *predicts 29260."*

Perhaps this comment dissuaded later researchers in matroid theory, because their catalogue remained unextended for more than 30 years, which surely makes it one of the longest standing computational results in combinatorics. However, in this paper we demonstrate that they were in fact unduly pessimistic, and describe an orderly algorithm (see McKay [7] and Royle [11]) that confirms their computations and extends them by determining the 383172 pairwise non-isomorphic matroids on nine elements (see Table 1).

Although this number of matroids is easily manageable on today's computers, our experiments with 10-element matroids suggests that there are at least 2.5×10^{12} sparse paving matroids of rank 5 on 10 elements. However we refrain from making any analogous predictions about the desirability or feasibility of constructing a catalogue of 10-element matroids!

We give some fundamental data about these matroids, and briefly describe how they are incorporated into an online database that provides access to a far greater range of data; this online database is accessible at `http://people.csse.uwa.edu.au/gordon/small-matroids.html`.

2 Matroids, flats and hyperplanes

We assume that the reader is familiar with the general idea of a matroid as being a combinatorial generalization of a multiset of vectors in a vector space, in that it consists of a set of elements and combinatorially defined concepts of dependence, independence etc. with properties analogous to the same concepts in vector spaces (see Oxley [8] for a gentle introduction to matroids, and Oxley [9] for complete details).

There are many equivalent ways to make this description precise, but for definiteness we will take a matroid $M = (E, r)$ to be a set E together with a rank function

$$r : 2^E \to \mathbb{Z}$$

satisfying the three conditions:

(R1) If $A \subseteq E$ then $0 \leq r(A) \leq |A|$.
(R2) If $A \subseteq B \subseteq E$ then $r(A) \leq r(B)$.
(R3) If $A, B \subseteq E$, then $r(A \cap B) + r(A \cup B) \leq r(A) + r(B)$.

The matroid equivalents of various concepts from linear algebra may then be defined using the rank function so that, for example, a set of elements $A \subseteq E$ is *independent* if $|A| = r(A)$, and *spanning* if $r(A) = r(E)$. The rank of the matroid itself is defined to be $r(E)$.

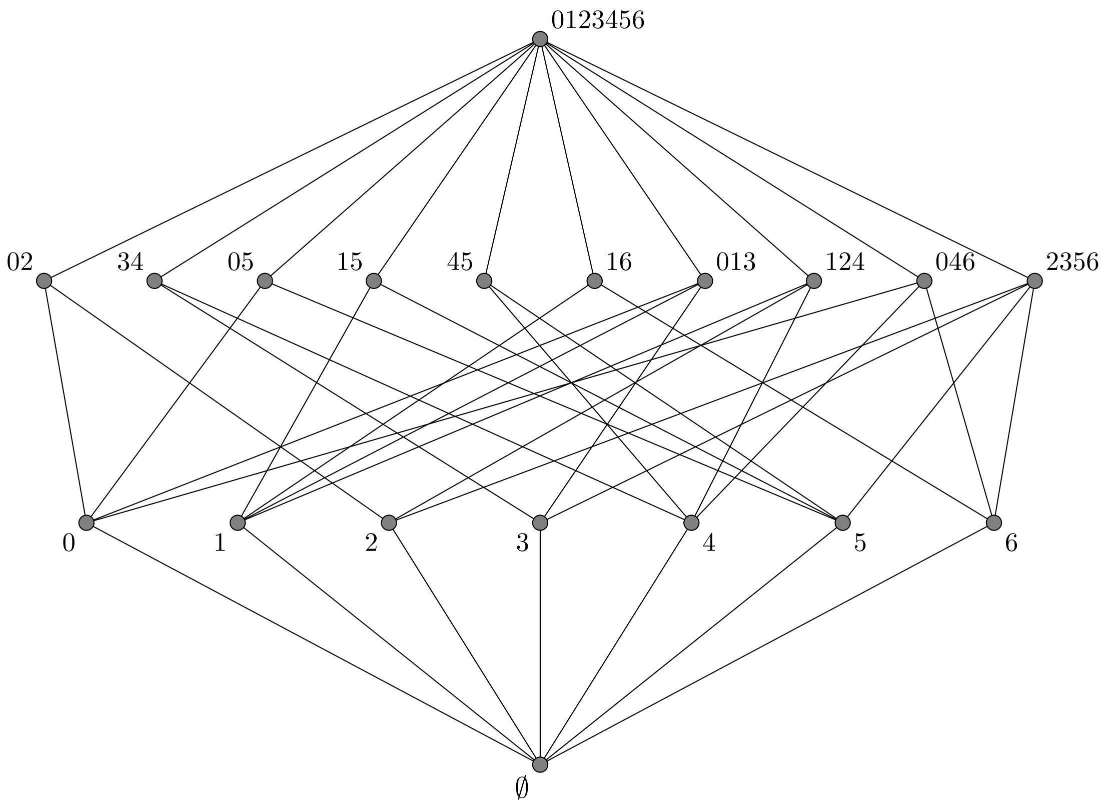

Figure 1: The lattice of flats of a rank 3 matroid on 7 elements

A *flat* is a set of elements $F \subseteq E$ such that for all $e \in E \backslash F$,

$$r(F \cup e) > r(F).$$

A flat is the matroid equivalent of a subspace, and as with subspaces, a crucial property of flats is that the intersection of two flats is again a flat. The intersection of all the flats containing a subset $A \subseteq E$ is called the *closure* of A and denoted by $\mathrm{cl}(A)$; it can be viewed as analogous to the subspace generated by a set of vectors.

The collection of flats of a matroid M under inclusion, together with the two operations

$$A \wedge B = A \cap B \qquad A \vee B = \mathrm{cl}(A \cup B)$$

forms a *geometric lattice* $\mathcal{L}(M)$. Figure 1 shows this lattice for a rank-3 matroid with 7 elements, where a label such as 013 is shorthand for the subset $\{0, 1, 3\}$. In this lattice we say that A *covers* B if $A \subset B$ and there are no flats C such that $A \subset C \subset B$. The rank function of M can be uniquely recovered from $\mathcal{L}(M)$ — the rank of a flat is the length of the maximal chain from $\emptyset$ to that flat, and the rank of any other subset $A \subseteq E$ is the rank of $\mathrm{cl}(A)$. In fact more is true, because the *hyperplanes* (i.e. the flats of rank $r(E) - 1$) are sufficient to determine the remaining flats, and hence the entire matroid.

Two matroids $M_1 = (E_1, r_1)$ and $M_2 = (E_2, r_2)$ are *isomorphic* if there is a bijection $\rho : E_1 \to E_2$ such that $r_2(\rho(A)) = r_1(A)$ for all $A \subseteq E$. For most (but not all) applications, it is appropriate to treat isomorphic matroids as equal and when counting and cataloguing matroids we are usually interested only in pairwise non-isomorphic matroids.

We can determine matroid isomorphism by using the *hyperplane graph* of the matroid, which is the bipartite graph whose vertices are the elements and hyperplanes of the matroid and where a hyperplane-vertex is adjacent to an element-vertex if and only if the hyperplane contains the element. From our previous discussion, it is clear that two matroids are isomorphic if and only if their hyperplane graphs are isomorphic as bipartite graphs (i.e. with the bipartition fixed). Although the theoretical complexity of graph isomorphism is not known, in practice Brendan McKay's program `nauty` [6] can easily process graphs with thousands of vertices (except for a few pathologically difficult, but poorly understood graphs). As the hyperplane graphs of the nine-element matroids have an average of only 74 vertices, isomorphism for matroids of this size is very easy to resolve in practice.

We note that using hyperplanes is a somewhat arbitrary choice and that any other collection of subsets that determines the matroid, such as the set of flats or the set of independent sets, could be used analogously.

3 Extensions and modular cuts

If $M = (E, r)$ is a matroid and $e \in E$, then the restriction of r to the subsets of $E\backslash e$ is itself a rank function, and so determines a matroid $M\backslash e = (E\backslash e, r\mid_{E\backslash e})$. We say that $M\backslash e$ is obtained by *deleting* e from M and conversely that M is an *single-element extension* of $M\backslash e$.

Now suppose that we have a list $\mathcal{M}_k$ of the matroids on k elements (or more precisely, one representative from each isomorphism class of matroids on k elements). Then we can form the list $\mathcal{M}_{k+1}$ of all matroids on $k+1$ elements by first finding all possible single-element extensions of every matroid in $\mathcal{M}_k$ and then eliminating unwanted isomorphic copies.

The key to extending a matroid in all possible ways lies in understanding the relationship between the flats of a matroid M and the flats of a single-element deletion $N = M\backslash e$. Letting $\mathcal{F}(M)$ denote the set of flats of a matroid M, it is easy to see that

$$\mathcal{F}(M\backslash e) = \{F\backslash e \mid F \in \mathcal{F}(M)\}. \tag{1}$$

Thus suppose that we are given the matroid N and wish to add a new element e, thereby finding all matroids M such that $M\backslash e = N$. By (1), every flat of M is of the form F or $F \cup \{e\}$ where $F \in \mathcal{F}(N)$. More precisely, for each flat $F \in \mathcal{F}(N)$ exactly one of the following three situations must hold in M:

1. $F \in \mathcal{F}(M)$ and $F \cup \{e\} \in \mathcal{F}(M)$
2. $F \in \mathcal{F}(M)$ but $F \cup \{e\} \notin \mathcal{F}(M)$
3. $F \notin \mathcal{F}(M)$ but $F \cup \{e\} \in \mathcal{F}(M)$

Thus the flats of $M \backslash e$ are partitioned into three parts in such a way that M can be uniquely recovered from $M \backslash e$ and this partition. Thus we can construct every possible single-element extension of a matroid N by considering all "suitable" partitions of $\mathcal{F}(N)$ into three parts and forming the different candidates for M accordingly.

This is feasible in practice because Crapo [4] showed that only certain highly structured partitions of $\mathcal{F}(M \backslash e)$ can actually arise, and therefore only a very limited number of partitions need be considered when extending a matroid. To describe this result we need one more piece of terminology: two flats $F, G \in \mathcal{F}(M)$ are a *modular pair* if

$$r(F) + r(G) = r(F \cup G) + r(F \cap G).$$

What Crapo showed was that if $N = M \backslash e$ and $\mathcal{F}(N) = \mathcal{F}_1 \cup \mathcal{F}_2 \cup \mathcal{F}_3$ is the partition of the flats of N according to the three possibilities listed above (respectively), then

1. $\mathcal{F}_3$ is an up-set in the lattice $\mathcal{L}(N)$ i.e. if $F \in \mathcal{F}_3$ then any flat containing F is in $\mathcal{F}_3$.
2. $\mathcal{F}_3$ is closed under taking intersections of *modular pairs* of flats.
3. $\mathcal{F}_2$ is the set of flats covered in $\mathcal{L}(N)$ by a member of $\mathcal{F}_3$.

The set $\mathcal{F}_3$ is called a *modular cut* and $\mathcal{F}_2$ the *collar* of the modular cut. As the modular cut determines its collar and $\mathcal{F}_1$ consists of the remaining flats, it follows that the modular cut alone determines the entire partition. Therefore we have the following result:

Theorem 1. *There is a 1-1 correspondence between modular cuts of N and single-element extensions of N.*

Figure 2 shows the modular cut $\{45, 0123456\}$ and the corresponding partition for the matroid of Figure 1.

The minimal elements of a modular cut form an anti-chain in $\mathcal{L}(N)$, and thus an easy way to determine the modular cuts of N is simply to compute all the anti-chains of $\mathcal{L}(N)$, form their up-sets and then check that the resulting set of flats is closed under intersection of modular pairs.

Blackburn, Crapo & Higgs used a more complicated scheme for computing modular cuts that avoids creating up-sets that are *not* modular cuts, but the overhead of the simpler scheme

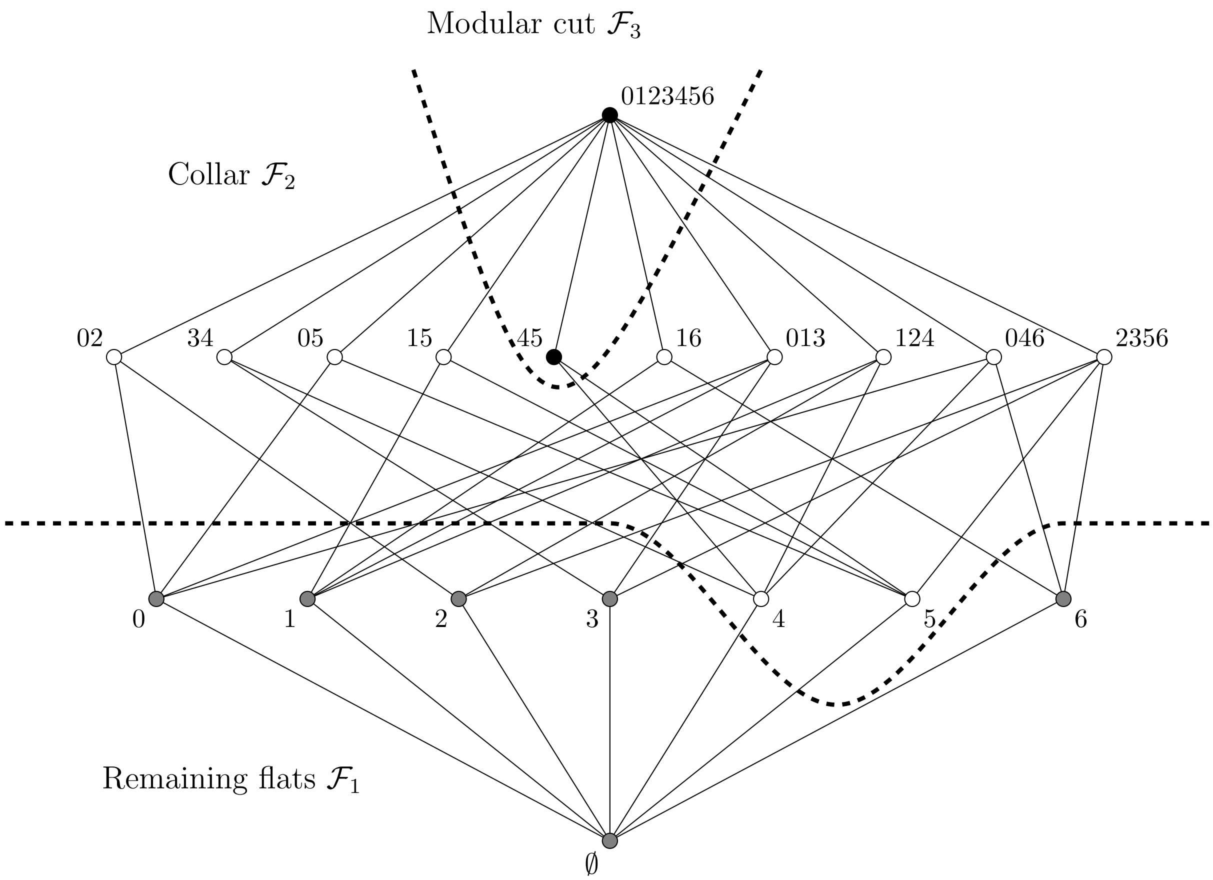

Figure 2: A modular cut (black nodes) and its collar (white nodes)

was sufficiently modest that we never had any need to implement the more complicated one. This also enhances our confidence in the correctness of our results in that existing very-well tested programs (for independent sets in graphs) could be used for computing anti-chains rather than necessarily less-tested bespoke programs.

4 An orderly algorithm

Our sole remaining task therefore is to consider matroid isomorphism and how to eliminate unwanted isomorphic copies of the matroids that are constructed, and for this we implemented a straightforward (partially) orderly algorithm (Read [10], McKay [7], Royle [11]).

In combinatorial construction, an *orderly algorithm* is one that is structured in such a way that it never outputs more than one representative of each isomorphism class of the objects being constructed, in this case matroids. What this means in practice is that as each matroid is produced by the extension procedure, it can be subjected to a test *not involving any other* *matroids* that determines whether it should be added to the output or rejected. Thus there is never any need to compare *pairs* of matroids, or test a newly-constructed matroid against a list of previously-constructed ones to check if it is really new.

Our algorithm falls into the category of "canonical construction path" orderly algorithms. Suppose that M is a matroid and that it has hyperplane graph $\mathcal{H}(M)$. Then `nauty` can be used to compute the canonical labelling of $\mathcal{H}(M)$ and thereby identify a distinguished element of M — for example, the element that receives the lowest canonical label. This then identifies a distinguished single-element deletion of M, namely the matroid obtained by deleting the distinguished element. The essence of the canonical construction path orderly algorithm is that it only accepts matroids that are constructed as an extension of this distinguished single-element deletion — whenever an isomorphic copy of M arises as an extension of one its *other* single-element deletions, it is rejected.

Although this ensures that the matroids generated by extending one matroid never need be compared with the matroids generated by extending a different matroid, it is still possible that two extensions from the *same* matroid may be isomorphic. Indeed this will necessarily happen if a matroid has two different, but isomorphic, modular cuts. However rather than perform isomorph rejection directly on modular cuts (many of which may lead to matroids that are subsequently rejected) we instead implemented simple "compare-and-filter" isomorph rejection on the set of matroids that were *accepted* when extending a single matroid.

Putting all this together, we get the procedure described in Algorithm 1.

Algorithm 1 Isomorph-free extension of a set X_k of k-element matroids

FOR EACH matroid $N \in X_k$ DO
    Set $N^+ \leftarrow \emptyset$.
    FOR EACH modular cut of N DO
        Form the single element extension M determined by the modular cut.
        Canonically label $\mathcal{H}(M)$ and add M to N^+ if and only if the newly added element is in the same orbit as the lowest canonically labelled element-vertex of $\mathcal{H}(M)$.
    END FOR
    Filter isomorphic matroids from N^+ and add the remainder to X_{k+1}.
END FOR

Notice that each matroid in X_k is processed entirely independently of the remaining matroids in X_k and therefore the computation can be arbitrarily partitioned between as many computers as desired.

Theorem 2. *If X_k contains one representative from each isomorphism class of k-element matroids, then X_{k+1} contains one representative from each isomorphism class of $(k+1)$-element matroids.*

Proof. Let M be an arbitrary $(k+1)$-element matroid, and let M' be its distinguished single-element deletion. By the hypothesis that X_k contains one representative from each isomorphism class of k-element matroids, a matroid isomorphic to M' will be processed at some stage, and so a matroid isomorphic to M will be constructed and then accepted. The filtering stage ensures that only one isomorph of M will be accepted during the processing of M' and the orderly aspect of the algorithm ensures that any isomorph of M is rejected whenever it is constructed as an extension of any matroid other than M'. Therefore X_{k+1} contains exactly one matroid isomorphic to M. □

5 Results

We implemented the algorithm described in the previous section, and the resulting numbers of matroids constructed are summarized in Table 1 (the totals form sequence A055545 in Neil Sloane's OEIS [12]).

$r\backslash n$	0	1	2	3	4	5	6	7	8	9
0	1	1	1	1	1	1	1	1	1	1
1		1	2	3	4	5	6	7	8	9
2			1	3	7	13	23	37	58	87
3				1	4	13	38	108	325	1275
4					1	5	23	108	940	190214
5						1	6	37	325	190214
6							1	7	58	1275
7								1	8	87
8									1	9
9										1
Total	1	2	4	8	17	38	98	306	1724	383172

Table 1: All matroids on up to 9 elements

These numbers are symmetric with respect to rank because of the theory of matroid *duality* — if $M = (E, r)$ is a matroid of rank r, then the function $r^* : 2^E \to \mathbb{Z}$ given by

$$r^*(A) = |A| + r(E\backslash A) - r(E)$$

is a rank function that determines a matroid $M^* = (E, r^*)$ of rank $|E|-r$ known as the *dual* of M. We emphasize that our algorithm did *not* exploit duality to reduce computation time by only constructing matroids of rank up to $|E|/2$, but rather we used the fact that our collection was closed under duality as a "sanity check" on the correctness of our implementation.

A *loop* in a matroid is an element of rank 0. If a matroid M has a loop ℓ then it appears in every flat, and the lattice $\mathcal{L}(M)$ has exactly the same structure as the lattice $\mathcal{L}(M\backslash\ell)$. Therefore loops play little structural role in M and it is common to consider them as trivial, and to consider only loopless matroids.

Two elements e and f are said to be *parallel* in a loopless matroid M if $r(\{e,f\}) = 1$. In this situation e and f are essentially duplicate elements and each flat either contains both of them or neither of them. Replacing $\{e,f\}$ with a single element wherever they occur does not alter the structure of the lattice $\mathcal{L}(\mathcal{M})$, and so again it is fairly common to consider them as trivial.

A matroid is called *simple* if it contains no loops and no parallel elements. Every matroid can be obtained from a unique simple matroid by adding a number of loops and parallel elements, and so in some sense the simple matroids are the "building blocks" from which all matroids can be constructed. Conversely, the unique simple matroid obtained from M by removing all loops and replacing each parallel class (i.e. set of mutually parallel elements) by a single element is called the *simplification* of M.

The catalogue of Blackburn, Crapo & Higgs only contains the *simple* matroids, and so we give their numbers in Table 2 and note that our computations are in complete agreement with theirs. In addition, Acketa [1] used Blackburn, Crapo & Higgs' catalogue to compute the numbers of all matroids (by adding loops and parallel elements in all possible ways) and our computations are also in complete agreement with his. More recently, Dukes [5] has given additional data about the matroids on up to 8 elements and again our results are in accordance with his.

$r\backslash n$	0	1	2	3	4	5	6	7	8	9
0	1									
1		1								
2			1	1	1	1	1	1	1	1
3				1	2	4	9	23	68	383
4					1	3	11	49	617	185981
5						1	4	22	217	188936
6							1	5	40	1092
7								1	6	66
8									1	7
9										1
Total	1	1	1	2	4	9	26	101	950	376467

Table 2: Simple matroids on up to 9 elements

We remark that declaring loops and parallel elements — but not their duals — to be trivial displays a somewhat *graph-theoretical* bias. In a matroid arising from a graph, a loop comes from a loop in the graph and parallel elements come from multiple edges, both of which are routinely excluded in much of graph theory. However, from a matroidal perspective, a matroid and its dual have equal status, and thus a loop is no more or less trivial than its dual, which is a *coloop*. Similarly, parallel elements are no more or less trivial than the dual structure, which are elements in *series*. In graphs, coloops correspond to cut-edges and elements in series correspond to paths with internal vertices of degree two. Graph theorists are understandably reluctant to declare these structures trivial because it would mean doing away with both trees and cycles! Matroidally however, the natural building blocks are those matroids that are both *simple* and *cosimple* (i.e. the dual matroid is also simple) and so we give their numbers in Table 3.

$r\backslash n$	0	1	2	3	4	5	6	7	8	9
0	1									
1										
2					1	1	1	1	1	1
3						1	6	20	65	380
4							1	20	525	185620
5								1	65	185620
6									1	380
7										1
Total	1	0	0	0	1	2	8	42	657	372002

Table 3: Simple and cosimple matroids on up to 9 elements

6 Paving Matroids

A *circuit* in a matroid is a minimal dependent set. It is possible for a matroid to have no circuits (in which case it consists entirely of coloops) but otherwise a matroid of rank r must have a circuit of size at most $r+1$. If the minimum circuit size is equal to $r+1$, then the matroid is a *uniform matroid* $U_{r,n}$ which has the property that the rank of a set A is equal to $\min(r, |A|)$. If the minimum circuit size is at least r, then the matroid is called a *paving matroid.*

We need some more terminology before we can understand why paving matroids form an important class of matroids.

A *d-partition* of a set E is a set $\mathcal{S}$ of subsets of E all of size at least d, such that every d-subset of E lies in a unique member of $\mathcal{S}$. Therefore a 1-partition of a set is simply a normal partition, while a 2-partition of a set is known as a *pairwise balanced design* with index 1. Obviously the set $\mathcal{S} = \{E\}$ is a d-partition for any d, and we call this the *trivial* d-partition.

The connection between paving matroids and d-partitions is given by the following result:

Theorem 3. *If $M = (E, r)$ is a paving matroid of rank $d+1 \geq 2$ then its hyperplanes form a non-trivial d-partition of E. Conversely, the elements of any non-trivial d-partition of E form the set of hyperplanes of a paving matroid of rank $d+1$.*

The *discrete* d-partition of a set E consists of all the d-subsets of E and the corresponding paving matroid is the uniform matroid $U_{d+1,|E|}$.

Based on the rather limited evidence in the catalogue of matroids on up to 8 elements, Welsh [14] asked whether *most* matroids are paving matroids. Examining the catalogue of 9-element matroids and tabulating the results in Table 4 we see that 71.71% of the simple matroids on 9 elements are paving matroids, compared to 49.50% of the 8-element simple matroids, thus providing some additional evidence that paving matroids do indeed predominate.

$r\backslash n$	0	1	2	3	4	5	6	7	8	9
0	1									
1		1								
2			1	1	1	1	1	1	1	1
3				1	2	4	9	23	68	383
4					1	2	5	18	322	147163
5						1	2	5	39	119050
6							1	2	6	178
7								1	2	6
8									1	2
9										1
Total	1	1	1	2	4	8	18	50	439	266784

Table 4: Simple paving matroids on up to 9 elements

A d-partition is called *sparse* if it contains no subsets of size greater than $d+1$, and similarly we call a paving matroid of rank $d+1$ sparse if its hyperplanes all have size d or $d+1$. A sparse paving matroid is determined completely by its hyperplanes of size $d+1$ — the d-partition must consist of these hyperplanes together with every d-set not yet contained in one of these. These hyperplanes are necessarily circuits, and so they form the set of *circuit-hyperplanes* of the matroid.

Sparse paving matroids have the attractive property that their duals are also sparse paving matroids — in fact the circuit-hyperplanes of M^* are the complements of the circuit-hyperplanes of M. Moreover if M and its dual are both paving matroids, then they are necessarily sparse and so the sparse paving matroids forms the largest possible dual-closed family of paving matroids.

Computationally, sparse paving matroids are attractive because they can be viewed simply as independent sets in a certain graph. The *Johnson graph* $J(n, d+1)$ is the graph whose vertices are all the $(d+1)$-subsets of an n-set, and where two vertices are adjacent if and only if the intersection of the corresponding subsets has size d. Therefore an *independent set* of vertices in $J(n, d+1)$ is precisely the set of circuit-hyperplanes of a sparse paving matroid of rank $d+1$, and conversely. Moreover the automorphism group of $J(n, d+1)$ is equal to the symmetric group S_n except in the special case where $n = 2(d+1)$ in which case the graph has an additional automorphism of order 2 induced by complementation on $(d+1)$-sets.

We note in passing that the *size* of the maximum independent set in the Johnson graphs

$J(n, d+1)$ has been intensively studied because such an independent set is directly equivalent to a *constant weight code* of length n, weight $d+1$ and minimum distance 4.

More generally, we can form the analogous graph on the $d+1$, $d+2$, …, $n-1$ sets of an n-set where again two vertices are adjacent if the corresponding sets meet in set of size d. Then an independent set in this graph corresponds to a not-necessarily-sparse paving matroid

7 Representability

A matroid $M = (E, r)$ of rank k is *representable* over a field $\mathbb{F}$ if there is a mapping

$$\rho : E \to \mathbb{F}^k$$

such that for any set $A \subseteq E$,

$$r(A) = \dim \mathrm{span}(\rho(A)).$$

To prove that a matroid *is* representable, it suffices to provide a suitable representation ρ, but it is considerably harder to prove that a matroid is *not* representable. However, Ingleton showed that if $M = (E, r)$ is representable and A, B, C, $D \subseteq E$, then

$$r(A) + r(B) + r(A \cup B \cup C) + r(A \cup B \cup D) + r(C \cup D) \\ \leq r(A \cup B) + r(A \cup C) + r(A \cup D) + r(B \cup C) + r(B \cup D).$$

It is therefore sometimes possible to show that a matroid is not representable by displaying four subsets A, B, C and D for which this inequality is violated. For want of a convenient term, we will call such matroids *Ingleton non-representable*.

Figure 3 gives a schematic diagram of the Ingleton-non-representable matroids on 8 elements; all of the matroids are sparse paving matroids and the diagram shows how they are related to each other under *relaxation of a circuit-hyperplane*, so that for example, the matroid F_8 is obtained from $AG(3, 2)'$ by relaxing a circuit-hyperplane. Any named matroids are listed according to their names in Oxley [9] while the remainder are given just by their number in the database. Dual pairs of matroids are connected by dotted lines.

In addition to the 39 Ingleton-non-representable matroids, there are five other rank-4 matroids on 8 elements that are non-representable. Four of these are related to the sparse paving matroid P_8. The matroid P_8 is a ternary matroid which has the representation

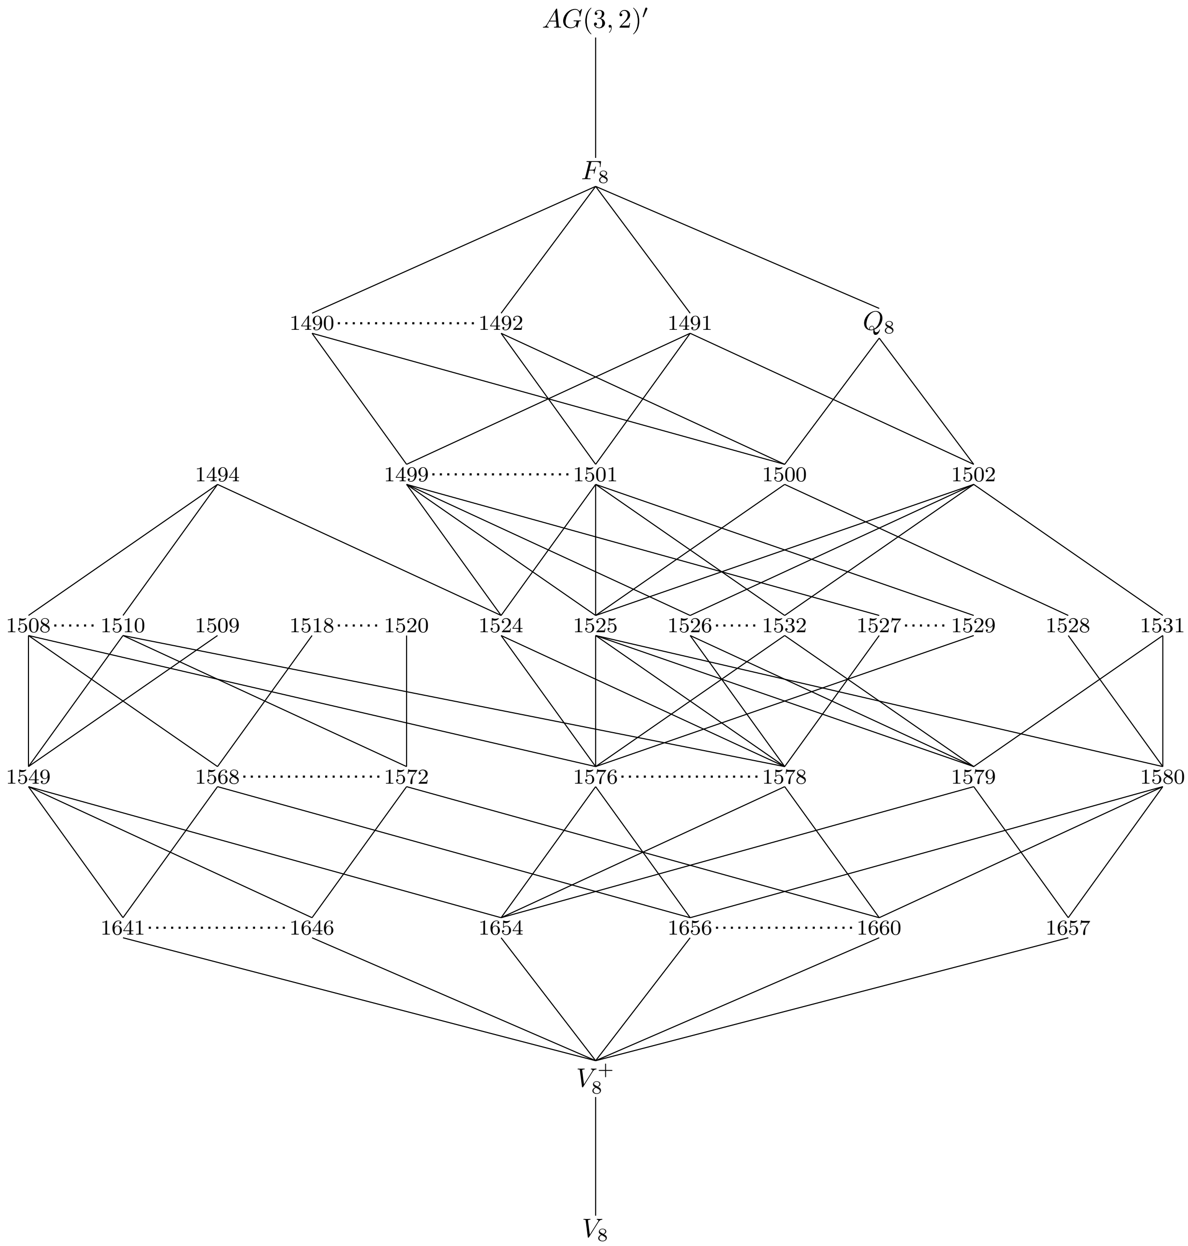

Figure 3: The 39 Ingleton-non-representable matroids on 8 elements

$$\begin{array}{cccccccc} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \end{array}$$
$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 1 & 1 & -1 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 & 1 \\ 0 & 0 & 1 & 0 & 1 & 1 & 0 & 1 \\ 0 & 0 & 0 & 1 & -1 & 1 & 1 & 0 \end{pmatrix}$$

The circuit-hyperplanes of this matroid are $\{0,1,2,3\}$, $\{0,1,3,6\}$, $\{0,2,3,5\}$, $\{1,2,3,4\}$, $\{0,3,4,7\}$, $\{1,2,5,6\}$, $\{0,4,5,6\}$, $\{1,4,5,7\}$, $\{2,4,6,7\}$ and $\{3,5,6,7\}$. We define four associated sparse paving matroids as follows: P_1 is obtained from P_8 by relaxing the circuit-hyperplane $\{3,5,6,7\}$, P_2' is obtained from P_1 by relaxing $\{0,3,4,7\}$, P_2'' is obtained from P_1 by relaxing $\{1,2,5,6\}$ and P_3 is obtained from P_1 by relaxing both $\{0,3,4,7\}$ and $\{1,2,5,6\}$.

Proposition 1. *The four matroids P_1, P_2', P_2'' and P_3 are all non-representable, but not Ingleton non-representable.*

Proof. Let $M \in \{P_1, P_2', P_2'', P_3\}$ and consider the basis $B = \{0,1,2,3\}$. Then following Section 6.4 of Oxley [9] a representation for M may be assumed to have the following form where $a,b,c,d,e \neq 0$ are unknown elements of some field.

$$A = \begin{array}{c} \begin{array}{cccccccc} 0 & 1 & 2 & 3 & 4 & 5 & 6 & 7 \end{array} \\ \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 1 & 0 & 0 & 1 & 0 & 1 & a \\ 0 & 0 & 1 & 0 & 1 & b & 0 & c \\ 0 & 0 & 0 & 1 & 1 & d & e & 0 \end{pmatrix} \end{array}$$

Now the sets $\{0,4,5,6\}$, $\{1,4,5,7\}$ and $\{2,4,6,7\}$ are circuits in M and so the submatrices of A defined on those particular sets of columns each have determinant 0. This gives the following three conditions respectively: $b(e-1)+d=0$, $b-c-d=0$ and $a+e-1=0$ which implies that

$$a = (1-e), \quad c = be \quad \text{and} \quad d = b(1-e)$$

However consider the submatrix of A with columns $\{3,5,6,7\}$. The determinant of this is

$$\begin{vmatrix} 0 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1-e \\ 0 & b & 0 & be \\ 1 & b(1-e) & e & 0 \end{vmatrix} = 0$$

contradicting the fact that $\{3,5,6,7\}$ is independent in M.

Checking Ingleton non-representability of a matroid is a task best left to a computer. □

The fifth non-representable matroid of size 8 for which Ingleton's condition gives no information is obtained from the matroid L_8 by relaxing a circuit-hyperplane, where L_8 is the sparse paving matroid whose circuit-hyperplanes are the 6 faces and the 2 colour-classes of a cube (see Oxley [9] p510). It can be shown to be non-representable using an analogous argument.

How effective is Ingleton's criterion for detecting non-representability among 9-element matroids? Perhaps surprisingly, it gives *no additional information* at all — that is, a 9-element matroid is Ingleton non-representable if and only if it contains an Ingleton-non-representable matroid on 8 elements as a minor.

8 A matroid database

One of the major uses of any sort of combinatorial catalogue is to compile data regarding the various combinatorial properties of the objects in the catalogue, and then to use this to answer questions or explore conjectures concerning the existence, or number of objects with various combinations of properties.

A common limitation of combinatorial catalogues is that their use is often restricted to their immediate creator and/or those researchers willing and able to download the raw data files and write their own programs, often resulting in significant duplication of effort. We have attempted to ameliorate this problem by incorporating the data into a relational database (using MySQL, although this is not important) and providing an online interface that permits "end users" to search, browse and investigate the data.

Currently we have computed a fairly substantial subset of what might be termed the "fundamental properties" of matroids. This includes various counts associated with each matroid such as numbers of loops, coloops, circuits, cocircuits, independent sets, bases, hyperplanes, flats and circuit-hyperplanes. It includes numerical properties such as the size of the automorphism group, the number of orbits of the automorphism group, the connectivity, the minimum circuit size. Various structural properties such as whether the matroid is binary, ternary, regular, paving, base-orderable, transversal and so on have also been included. More importantly however, we have incorporated information about the *relationships* between the matroids — relationships such as duality, deletion and contraction of elements, relaxation of circuit-hyperplanes, truncations and simplifications. Finally we have included auxiliary information such as information about rank polynomials and representations over small finite fields.

Rather than present a large number of tables of data in this paper, we give three simple examples of the use of the database, and invite readers to explore their own particular interests by using the database at `http://people.csse.uwa.edu.au/gordon/small-matroids.html`.

8.1 Excluded minors for $GF(5)$

One of the most fundamental results in matroid theory is Tutte's characterization of matroids representable over $GF(2)$ in terms of excluded minors: a matroid is binary if and only if it does not contain $U_{2,4}$ as a minor.

Similar characterizations are known for matroids representable over $GF(3)$ where there are 4 excluded minors and $GF(4)$ where there are 7 excluded minors. The analogous characterization for $GF(5)$ is not known or even conjectured, with prevailing opinion suggesting that such a characterization is likely to be extremely complex and unwieldy. In fact, Whittle [15] suggests that "It is not clear that the problem for finding the specific excluded minors for $GF(5)$ is that well motivated" and that the real question in representability is to resolve Rota's conjecture that the list of excluded minors for representability over any field is finite.

It is straightforward to determine the matroids on up to 9 elements that *are* representable over $GF(5)$ by finding all sets of at most 9 points in the projective space $PG(3,5)$ that are pairwise inequivalent under the action of the group $PGL(4,5)$, identifying the corresponding matroids (which have rank at most 4) and then finding their duals.

Once this is done, we can identify the matroids that are *not* $GF(5)$-representable but for which every single-element deletion and single-element contraction *is* $GF(5)$-representable and thus determine the excluded minors on at most 9 elements. Table 5 shows the numbers of matroids that were found, confirming the belief that an excluded minor characterization of $GF(5)$-representable matroids in the traditional style would indeed be very cumbersome.

Size	Rank	No.	Comment
7	2	1	Uniform $U_{2,7}$
7	3	5	
7	4	5	
7	5	1	Uniform $U_{5,7}$
8	3	2	
8	4	92	
8	5	2	
9	3	9	
9	4	219	
9	5	219	
9	6	9	

Table 5: Excluded minors for $GF(5)$ on up to nine elements

8.2 Numbers of bases

In a matroid of rank r on n elements, the number b of bases must necessarily satisfy $1 \leq b \leq \binom{n}{r}$. In 1969, Welsh [13] conjectured that for *every triple* (n, r, b) such that $0 \leq r \leq n$ and $1 \leq b \leq \binom{n}{r}$, there is a matroid of rank r on n elements with exactly b bases — in other words, everything that can happen, does.

We can check this all matroids on up to 9 elements with a single SQL statement (though note that `binomial` is not a built-in function, but must be programmed):

```
SELECT tmp.size, tmp.rank, COUNT(*) FROM
  (SELECT DISTINCT size, rank, numBases FROM matroids9) as tmp
  GROUP BY tmp.size, tmp.rank
  HAVING COUNT(*) <> binomial(tmp.size, tmp.rank);
```

The inner `SELECT` statement first creates a list of all the distinct triples (n, r, b) represented in the database and gives it the alias `tmp`. The outer `SELECT .... GROUP BY` statement *counts* the triples in `tmp` for each fixed pair (n, r), while the `HAVING` statement extracts the pairs where this count is not equal to $\binom{n}{r}$, thus representing one or more "missing" triples.

```
+------+------+----------+
| size | rank | COUNT(*) |
+------+------+----------+
|    6 |    3 |       19 |
+------+------+----------+
```

The output shows that there are only 19 triples of the form $(6, 3, b)$, rather than the expected 20. In fact, the missing triple is $(6, 3, 11)$ — there are no rank-3 matroids on 6 elements with exactly 11 bases — a fact that was previously observed by Anna de Mier (personal communication). The absence of any other missing triples with $n \leq 9$ and the exponential explosion in numbers of matroids as n reaches 10 leads us to strongly believe the following conjecture:

Conjecture 1. *For every triple* (n, r, b) *such that* $0 \leq r \leq n$ *and* $1 \leq b \leq \binom{n}{r}$ *there is a matroid of rank* r *on* n *elements with exactly* b *bases* except when

$$(n, r, b) = (6, 3, 11).$$

8.3 Transversal matroids

Given two bases A and B of a matroid, the subsets $X \subseteq A$ and $Y \subseteq B$ are called *exchangeable* if both $A \backslash X \cup Y$ and $B \backslash Y \cup X$ are bases.

Rank \Size	2	3	4	5	6	7	8	9
2	1	3	7	13	23	37	58	87
	1	3	7	13	23	37	58	87
	1	3	7	13	23	37	58	87
	1	3	7	13	22	34	50	70
3		1	4	13	38	108	325	1275
		1	4	13	37	101	284	956
		1	4	13	37	101	284	956
		1	4	13	37	92	209	442
4			1	5	23	108	940	190214
			1	5	23	101	677	70569
			1	5	23	101	644	55081
			1	5	23	100	432	1804
5				1	6	37	325	190214
				1	6	37	284	70569
				1	6	37	284	55081
				1	6	37	272	2806
6					1	7	58	1275
					1	7	58	956
					1	7	58	956
					1	7	58	817

Table 6: Numbers of matroids, base-orderable matroids, strongly base-orderable matroids and transversal matroids.

A matroid is called *base-orderable* if for any two bases A and B there is a bijection $\varphi : A \to B$ such that a and $\varphi(a)$ are exchangeable, while it is *strongly base-orderable* if the bijection can be selected so that X and $\varphi(X)$ are exchangeable for every *subset* $X \subseteq A$. A matroid on a set E is a *transversal matroid* if there is a bipartite graph G with bipartition $E \cup F$ such that the independent sets of M are precisely the subsets of E that are the endpoints of a matching (i.e. an independent set of edges) of G.

It is obvious that strongly base-orderable matroids are base-orderable, but less obvious that transversal matroids and their duals (which need not be transversal matroids) are strongly-base orderable. These classes of matroids are important but not fully understood, and therefore the numbers of matroids in each of these classes is of some interest. Determining whether a matroid is base-orderable or strongly base-orderable is straightforward, and we implemented the algorithm given by Brualdi & Dinolt [3] for testing transversality.

Table 6 gives these numbers where each cell of the table contains four numbers which, reading from top to bottom are the total number of matroids and the number of base-orderable, strongly base-orderable and transversal matroids respectively. For the omitted ranks (ranks 0, 1, 7, 8 and 9) all the matroids in the catalogue are transversal.

Size	Number	Size	Number	Size	Number	Size	Number
0	1	8	521367	16	579539500	24	1355
1	2	9	3539486	17	329728133	25	250
2	3	10	18146294	18	130254690	26	58
3	13	11	69516384	19	35087875	27	13
4	73	12	197898106	20	6400127	28	4
5	575	13	416277780	21	818999	29	1
6	5838	14	642315652	22	84722	30	1
7	59818	15	720126836	23	9263		

Table 7: Independent sets in $J(10,4)$

9 Matroids on ten elements?

Given that 30+ years have elapsed since the catalogue of matroids on 8 elements was created and with the benefit of advances both in raw computational power and techniques in combinatorial construction, it may seem rather unambitious to extend the catalogue only to 9 elements.

However our initial experiments on the feasibility of constructing the matroids on 10 elements lead us to the conclusion that even *counting* the 10-element matroids would be a major undertaking, let alone constructing them.

This is very unfortunate because we have a strong feeling that the absence of geometric representations available for rank-3 and rank-4 matroids means that rank-5 matroids are in some sense much less well understood than their lower rank counterparts. One of our original motivations in embarking on this project was the belief that the rank-5 matroids on 10 elements might be a fertile source of interesting and/or counterintuitive examples and counterexamples for this reason.

9.1 Paving matroids of rank 4

From our analysis above, the sparse paving matroids of rank 4 on 10 elements are in 1-1 correspondence with independent sets in the Johnson graph $J(10,4)$, with isomorphism of matroids and isomorphism under the automorphism group S_{10} of the graph being the same. Therefore a straightforward orderly algorithm as outlined in Royle [11] can be used to construct them. This computation was performed in a few days using idle time on a network of about 50 computers, and the resulting numbers are presented in Table 7 which shows a total of 3150333219 (i.e. $\approx 3.150 \times 10^9$) sparse paving matroids of rank 4 on 10 elements.

Computation of the *non-sparse* paving matroids of rank 4 on 10 elements is a somewhat

Max. hyp. k	No. k-hyps.	No. matroids
5	1	1222076172
5	2	147724716
5	3	5558695
5	4	64194
5	5	232
5	6	6
6	1	2369590
6	2	164
7	1	435
8	1	5
9	1	1
Total		1377794210

Table 8: Non-sparse paving matroids of rank 4 on 10 elements

fiddly bookkeeping exercise, but it involves no qualitatively different techniques. The essence of our approach is to divide the search according to whether the largest hyperplane has size $k = 5$, 6, 7, 8 or 9. For each size k, we construct an auxiliary graph $G(k)$ defined on the 4-, 5-, ..., k-sets that meet a fixed k-set (e.g. $\{0, 1, \ldots, k-1\}$) in less than 3 points and with adjacency again defined by intersection in at least 3 points. Then an independent set of $G(k)$ together with $\{0, 1, \ldots, k-1\}$ and all 3-sets not already covered forms the set of hyperplanes of a non-sparse paving matroid. However we need to be a little careful with isomorphism — this procedure distinguishes a particular k-set and so if a matroid has c orbits on hyperplanes of size k, then it will contribute c pairwise non-isomorphic independent sets to $G(k)$. Therefore each independent set contributes $1/c$ to the total count of matroids, where c is the number of orbits that the corresponding matroid has on hyperplanes of size k. Of course if the matroid has only one hyperplane of size k, then $c = 1$ follows immediately with no special calculation.

Table 8 shows the results of this calculation broken down according to the size of the largest hyperplane k and how many hyperplanes of this size are in the matroid.

Adding the numbers of sparse and non-sparse paving matroids, we conclude that there are 4528127429 ($\approx 4.528 \times 10^9$) paving matroids of rank 4 on 10 elements.

9.2 Paving matroids of rank 5

We have been unable to complete the analogous computation for the sparse paving matroids of rank 5 on 10 elements. These correspond to independent sets in the Johnson graph $J(10, 5)$ but with one additional complication. The automorphism group of $J(10, 5)$ is $S_{10} \times Z_2$ with the additional Z_2 being induced by complementation of 5-sets. This means that each independent set of $J(10,5)$ produced by the orderly algorithm corresponds to a dual pair of matroids — usually two matroids, but only one when the matroid is self-dual. Thus the total number of matroids is twice the number of independent sets of $J(10,5)$ minus the number of self-dual matroids.

We can determine the number of self-dual sparse paving matroids on 10 elements in a separate computation by exploiting the fact that the corresponding independent sets must have a non-trivial automorphism involving the Z_2 part of the automorphism group of $J(10,5)$. This separate computation yields a total of 99022169 self-dual sparse paving matroids.

However the sheer number of independent sets in $J(10,5)$ makes it infeasible for us to complete the first part of the computation. We can however make an "informed guess" of the magnitude of the number by executing a fixed percentage of the search. First, the orderly algorithm was used to compute the entire collection of independent sets of size 9, of which there are 20680075. A random sample of this collection was selected, and the search completed just using these as the starting points. Although it is hard to say anything statistically precise, our prior experience with such orderly algorithms suggest that the number of independent sets produced is roughly proportional to the size of the random sample of starting points.

A sample of 60000 starting points (0.2901% of the search space) yielded 3.875×10^9 independent sets giving an estimate of 1.336×10^{12} independent sets in $J(10,5)$. Therefore we estimate that there are about 2.65×10^{12} sparse paving matroids of rank 5 on 10 elements.

References

[1] Acketa, D. The catalogue of all nonisomorphic matroids on at most 8 elements, vol. 1 of Special Issue. University of Novi Sad Institute of Mathematics Faculty of Science, Novi Sad, 1983.

[2] Blackburn, J. E., Crapo, H. H., and Higgs, D. A. A catalogue of combinatorial geometries. Math. Comp. 27 (1973), 155–166; addendum, ibid. 27 (1973), no. 121, loose microfiche suppl. A12–G12.

[3] Brualdi, R. A., and Dinolt, G. W. Characterizations of transversal matroids and their presentations. J. Combinatorial Theory Ser. B 12 (1972), 268–286.

[4] Crapo, H. H. Single-element extensions of matroids. J. Res. Nat. Bur. Standards Sect. B 69B (1965), 55–65.

[5] Dukes, W. M. B. On the number of matroids on a finite set. Sém. Lothar. Combin. 51 (2004/05), Art. B51g, 12 pp. (electronic).

[6] McKay, B. D. Practical graph isomorphism. In Proceedings of the Tenth Manitoba Conference on Numerical Mathematics and Computing, Vol. I (Winnipeg, Man., 1980) (1981), vol. 30, pp. 45–87. Available from: http://cs.anu.edu.au/~bdm/nauty/PGI/.

[7] McKay, B. D. Isomorph-free exhaustive generation. J. Algorithms 26, 2 (1998), 306–324.

[8] Oxley, J. What is a matroid? Cubo Mat. Educ. 5, 3 (2003), 179–218.

[9] Oxley, J. G. Matroid theory. Oxford Science Publications. The Clarendon Press Oxford University Press, New York, 1992.

[10] Read, R. C. Every one a winner or how to avoid isomorphism search when cataloguing combinatorial configurations. Ann. Discrete Math. 2 (1978), 107–120. Algorithmic aspects of combinatorics (Conf., Vancouver Island, B.C., 1976).

[11] Royle, G. F. An orderly algorithm and some applications in finite geometry. Discrete Math. 185, 1-3 (1998), 105–115.

[12] Sloane, N. J. A. The On-Line Encylopaedia of Integer Sequences. Available from: http://www.research.att.com/~njas/sequences.

[13] Welsh, D. J. A. Combinatorial problems in matroid theory. In Combinatorial Mathematics and its Applications (Proc. Conf., Oxford, 1969). Academic Press, London, 1971, pp. 291–306.

[14] Welsh, D. J. A. Matroid theory. Academic Press [Harcourt Brace Jovanovich Publishers], London, 1976. L. M. S. Monographs, No. 8.

[15] Whittle, G. Recent work in matroid representation theory. Discrete Math. 302, 1-3 (2005), 285–296.