On Some Hard Problems on Matroid Spikes

Petr Hliněný[⋆]

Department of Computer Science,
VŠB – Technical University Ostrava,
17. listopadu 15, 708 33 Ostrava, Czech Republic
and
Faculty of Informatics,
Masaryk University in Brno,
Botanická 68a, 602 00 Brno, Czech Republic
hlineny@fi.muni.cz

November 30, 2005

Abstract. Spikes form an interesting class of 3-connected matroids of branch-width 3. We show that some computational problems are hard on spikes with given matrix representations over infinite fields. Namely, the question whether a given spike is the free spike is co-NP-hard (though the property itself is definable in monadic second-order logic); and the task to compute the Tutte polynomial of a spike is $\#P$-hard (even though that can be solved efficiently on all matroids of bounded branch-width which are represented over a finite field).

Keywords: matroid, branch-width, spike, computational complexity, Tutte polynomial.
2000 Math subject classification: 05B35, 68Q17, 68R05.

1 Introduction

We postpone necessary formal definitions until later sections. Among successful structural (combinatorial) parameters in the theory of parametrized complexity [2], the prominent role is played by tree-width, or equivalently, branch-width. Considering graphs, almost all usual hard problems can be efficiently solved on graphs of bounded tree-width. Those include all problems definable in the (monadic second-order) MSO logic of graphs, or the notoriously hard graph counting invariant, the Tutte polynomial.

It has recently turned out that many of those efficient parametrized results carry over also to matroids of bounded branch-width which are represented by matrices over finite fields. (A matrix representation of a matroid is the most common one studied.) However, one can observe a radical structural change when getting to matroid representations over infinite fields: For example [4], $\mathbb{F}$-representable matroids of bounded branch-width are well-quasi-ordered under

⋆ The research has been supported by Czech research grant GAČR 201/05/0050, and by the Institute of Theoretical Computer Science, project 1M0545.

the minor ordering for any finite field $\mathbb{F}$. On the other hand, an interesting class of matroids of branch-width three, called spikes, is not well-quasi-ordered over any infinite field.

Our research extends evidence for this structural difference of matroid representations further to algorithmic questions.

- By our results in [7, 6], all the matroid properties expressible in MSO logic are recognizable in cubic time for matroids of bounded branch-width which are represented over finite fields. On contrary, we show here a simple MSO formula which is NP-complete to decide on matroid spikes which are represented over any infinite field. (Theorem 3.2)
- Analogously, for counting problems, the Tutte polynomial can be efficiently computed [5] on matroids of bounded branch-width which are represented over finite fields. We prove here that this invariant is $\#P$-hard on matroid spikes which are represented over any infinite field. (Theorem 4.3)

Our proofs use a direct reduction from the PARTITION problem – one of the 6 basic NP-complete problems in [3].

2 Basics of Matroids

We refer to Oxley [14] in our matroid terminology. A *matroid* is a pair $M = (E, \mathcal{B})$ where $E = E(M)$ is the ground set of M (elements of M), and $\mathcal{B} \subseteq 2^E$ is a nonempty collection of *bases* of M, no two of which are in an inclusion. Moreover, matroid bases satisfy the "exchange axiom": if $B_1, B_2 \in \mathcal{B}$ and $x \in B_1 - B_2$, then there is $y \in B_2 - B_1$ such that $(B_1 - \{x\}) \cup \{y\} \in \mathcal{B}$. We consider only finite matroids. Subsets of bases are called *independent sets*, and the remaining sets are *dependent*. Minimal dependent sets are called *circuits*. All bases have the same cardinality called the *rank* $r(M)$ of the matroid. The *rank function* $r_M(X)$ in M is the maximal cardinality of an independent subset of a set $X \subseteq E(M)$.

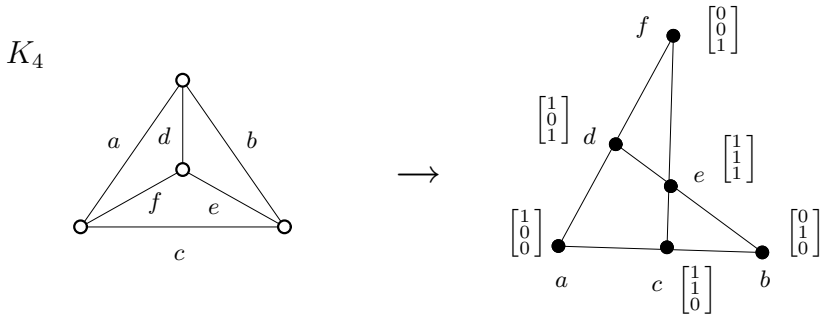

Fig. 1. An example of a vector representation of the cycle matroid $M(K_4)$. The matroid elements are depicted by dots, and their (linear) dependency is shown using lines.

If G is a (multi)graph, then its *cycle matroid* on the ground set $E(G)$ is denoted by $M(G)$. The independent sets of $M(G)$ are acyclic subsets (forests) in G, and the circuits of $M(G)$ are the cycles in G. Another example of a matroid is a finite set of vectors with usual linear dependency. If $\boldsymbol{A}$ is a matrix, then the matroid formed by the column vectors of $\boldsymbol{A}$ is called the *vector matroid* of $\boldsymbol{A}$, and denoted by $M(\boldsymbol{A})$. The matrix $\boldsymbol{A}$ is a *representation* of a matroid $M \simeq M(\boldsymbol{A})$. We say that the matroid $M(\boldsymbol{A})$ is $\mathbb{F}$*-represented* if $\boldsymbol{A}$ is a matrix over a field $\mathbb{F}$. (Fig. 1.) A *graphic matroid*, i.e. a cycle matroid of some multigraph, is representable over any field.

An interesting question about matroids arises in connection with computational complexity: What is the input size of an n-element matroid? In truth, it is $\Theta(2^n)$ since a matroid carries information about all subsets of its ground set, but acceptance of that would ruin usual algorithmic complexity measures. (Though, there is a recent work of Mayhew [12] on the complexity of problems on matroids given as exponential set-lists.) That is why matroids are often considered with particular representations of polynomial size, like the above mentioned graphic or vector matroids.

Matroid Branch-Width

Since the notion of branch-width is less known than that of tree-width, we briefly introduce it from a matroidal point of view. A *sub-cubic tree* is a tree in which all vertices have degree at most three. Let $\ell(T)$ denote the set of leaves of a tree T. The *connectivity* function λ_M of a matroid M on E is defined for all $A \subseteq E$ by $\lambda_M(A) = r_M(A) + r_M(E - A) - r_M(E) + 1$.

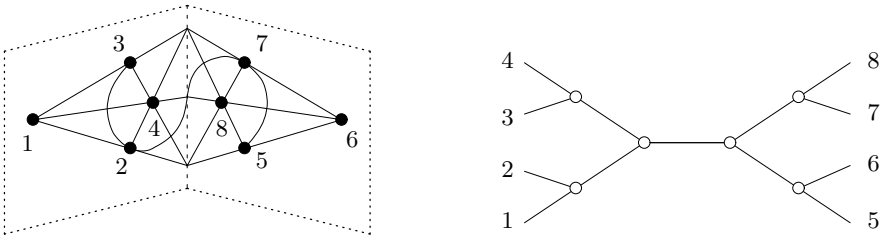

Fig. 2. A width-3 branch-decomposition of a matroid (the binary cube).

A *branch-decomposition* of a matroid M is a pair (T, τ) where T is a sub-cubic tree, and τ is a bijection of E onto $\ell(T)$. Let e be an edge of T, and T_1, T_2 be the connected components of $T - e$. The *width* of the edge e in T is $\lambda_M(F_e) = \lambda_M(E - F_e)$ where $F_e = \tau^{-1}(\ell(T_1))$. The width of the branch-decomposition (T, τ) is maximum of the widths of all edges of T, and the *branch-width* of M is the minimal width over all branch-decompositions of M. (This definition is similar to branch-width of graphs. See in Fig. 2.)

Introducing Spikes

We now introduce an interesting class of matroids, called "spikes". Let $n \geq 3$ and S_0 be a matroid circuit on the ground set $e_0, e_1, \ldots, e_n$. Denote by S_1 an arbitrary simple matroid obtained from S_0 by adding n new elements f_i for $i \in [1, n]$ such that $\{e_0, e_i, f_i\}$ is a triangle. Then the matroid $S = S_1 \setminus e_0$ obtained by deleting the central element e_0 is called a *rank-n spike*. The pairs $\{e_i, f_i\}$, $i \in [1, n]$ are called the *legs* of the spike. (Fig. 3.)

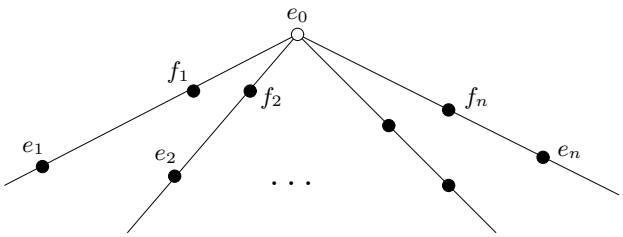

Fig. 3. An illustration to the definition of a rank-n spike.

Spikes are known for giving "difficult counterexamples", and they more or less explicitly appear in several papers in structural matroid theory, among recent we mention [4]. We remark that some of the findings in this section appear implicitly already in [15]. There seems to be no "usual definition" of a spike; the above definition was suggested by Whittle. The following basic properties of spikes are well known in the matroid community, and so we only sketch their proofs here.

Proposition 2.1. *Let S be a rank-n spike where $n \geq 3$. Then*

a) the union of any two legs forms a 4-element circuit in S,
b) every other circuit intersects all legs of S,
c) every set of elements that contains at most one leg and is disjoint from some other leg is independent, and
d) the branch-width of S is 3.

Sketch of proof. We use the notation (S_1, S) from the definition of a spike.

(a) The lines of any two legs are coplanar since they intersect in e_0, and no proper subset of those 4 elements is dependent.

(b) Let C be a circuit in S disjoint from the leg $\{e_1, f_1\}$, and containing no two legs. Then C contains a circuit of the minor $S' = S_1/e_0 \setminus \{e_1, f_1\}$. Since there are no other circuits in S' than parallel pairs of elements (former legs of S) by definition, it must be that C contained some leg, say $e_2, f_2 \in C$. However, now the same argument applies to the minor $S'' = S/\{e_2, f_2\} \setminus \{e_1, f_1\}$ (which is isomorphic to $S'/e_2 \setminus f_2$), and hence C contained another leg of S, a contradiction.

(c) This is immediate from (a) and (b) since each dependent set contains a circuit.

(d) Let us take an arbitrary cubic tree with n leaves, and attach to each leaf two new ones labelled with e_i and f_i, $i = 1, 2, \ldots, n$. It is routine to verify that the width of this decomposition is 3. ∎

In particular, the proposition completely describes all circuits (or dependent sets) of size less than n for $n \geq 5$. Representable spikes have particularly nice matrix representations, as shown in Fig. 4. Moreover, one may simply read off the matroid structure from such a representation. Let $\boldsymbol{D}^1(x_1, \ldots, x_n) = [d_{i,j}]_{i=1}^n$ denote an $n \times n$ matrix such that $d_{i,j} = 1$ if $i \neq j$, and $d_{i,i} = x_i$ for $i \in [1, n]$.

$$
\begin{array}{c}
\\
e_1 \\ e_2 \\ \vdots \\ e_{n-1} \\ e_n
\end{array}
\begin{array}{c}
\begin{array}{ccccccccccc}
e_1 & e_2 & \ldots & e_{n-1} & e_n & & f_1 & f_2 & \ldots & f_{n-1} & f_n
\end{array} \\
\left[\begin{array}{ccccccccccc}
1 & 0 & \cdots & 0 & 0 & & x_1 & 1 & \cdots & 1 & 1 \\
0 & 1 & 0 & 0 & 0 & & 1 & x_2 & 1 & 1 & 1 \\
\vdots & 0 & \ddots & 0 & \vdots & & \vdots & 1 & \ddots & 1 & \vdots \\
0 & 0 & 0 & 1 & 0 & & 1 & 1 & 1 & x_{n-1} & 1 \\
0 & 0 & \cdots & 0 & 1 & & 1 & 1 & \cdots & 1 & x_n
\end{array}\right]
\end{array}
\begin{array}{c}
\\
e_1 \\ e_2 \\ \vdots \\ e_{n-1} \\ e_n
\end{array}
$$

Fig. 4. A matrix representation of a rank-n spike ($x_i \neq 1$).

Lemma 2.2. *Let $\mathbb{F}$ be any field and $n \geq 3$. A rank-n spike S is representable over $\mathbb{F}$ if and only if S is represented by the matrix $\left[\boldsymbol{I}_n \,|\, \boldsymbol{D}^1(x_1, \ldots, x_n)\right]$ for some $x_1, \ldots, x_n \in \mathbb{F} - \{1\}$.*

Proof. Consider a spike S. We represent the basis $\{e_1, \ldots, e_n\}$ of S by the unit vectors of $\boldsymbol{I}_n$. Then the element e_0 from the definition of a spike is, up to scaling, represented by the vector $\boldsymbol{1}$. Hence the element f_i colinear with e_0, e_i is represented by a vector $\boldsymbol{1} + (x_i - 1)\boldsymbol{e}_i$ for $i \in [1, n]$ and some $x_i \in \mathbb{F} - \{1\}$. The proof of the converse is clear. ∎

Lemma 2.3. *Let $\boldsymbol{D}_k = \boldsymbol{D}^1(y_1, \ldots, y_k)$. If $y_i = 1$ for more than one index $i \in [1, k]$, then the determinant $|\boldsymbol{D}_k| = 0$. If $y_i = 1$ for exactly one index $i \in [1, k]$, then $|\boldsymbol{D}_k| \neq 0$. Otherwise, if $y_i \neq 1$ for all $i \in [1, k]$, then*

$$|\boldsymbol{D}_k| = \left|\boldsymbol{D}^1(y_1, \ldots, y_k)\right| = \left(\prod_{i=1}^{k}(y_i - 1)\right) \cdot \left(1 + \sum_{i=1}^{k} \frac{1}{y_i - 1}\right).$$

Proof. We use simple matrix row operations:

$$|\boldsymbol{D}_k| = \begin{vmatrix} y_1 & 1 & \cdots & 1 \\ 1 & y_2 & \cdots & 1 \\ \vdots & & \ddots & \vdots \\ 1 & 1 & \cdots & y_k \end{vmatrix} = \begin{vmatrix} y_1 & 1 & \cdots & 1 \\ 1 - y_1 & y_2 - 1 & \cdots & 0 \\ \vdots & & \ddots & \vdots \\ 1 - y_1 & 0 & \cdots & y_k - 1 \end{vmatrix}$$

The answer is clear if some $y_i = 1$. Otherwise, we expand the determinant as

$$|\boldsymbol{D}_k| = (y_1 - 1)\left(\left(\frac{1}{y_1 - 1} + 1\right)\prod_{i=2}^{k}(y_i - 1) + \sum_{i=2}^{k}\prod_{j=2,\, j\neq i}^{k}(y_j - 1)\right). \qquad \blacksquare$$

When dealing with matroids represented by matrices in the standard form $\boldsymbol{A} = [\boldsymbol{I} \mid \boldsymbol{A'}]$, it is crucial to notice that the matroid bases are in a one-to-one correspondence with the nonzero subdeterminants in $\boldsymbol{A'}$ (via standard matrix pivoting). Hence the following claim is just a reformulation of Proposition 2.1 (b,c), using the determinant formula of Lemma 2.3.

Corollary 2.4. *Let S be a rank-n spike represented over $\mathbb{F}$ by the matrix $\big[\boldsymbol{I}_n \mid \boldsymbol{D}^1(x_1, \ldots, x_n)\big]$, $n \geq 5$, and let $E(S) = \{e_1, \ldots, e_n, f_1, \ldots, f_n\}$ as in Fig. 4. Consider a set $X \subseteq E(S)$.*

a) If there are indices $i, i' \in [1, n]$ such that $e_i, f_i \in X$, $e_{i'}, f_{i'} \notin X$, and that $|\{e_j, f_j\} \cap X| = 1$ for all $j \neq i, i'$, then X is a basis of S.
b) If $|\{e_j, f_j\} \cap X| = 1$ for all indices $j \in [1, n]$, then X is a basis of S if and only if

$$\sum_{j \in [1,n],\, f_j \in X} \frac{1}{x_j - 1} \neq -1\,.$$

c) Otherwise, X is not a basis of S.

$\blacksquare$

3 Easy and Hard MSO Properties

Monadic second-order (MSO) logic is famous for its well-balanced expressive power and algorithmic manageability of its theories. Namely, considering MS_2-definable properties of graphs (quantifying over vertex and edge sets), all those can be efficiently solved on graphs of bounded tree-width [1]. An analogous phenomenon appears for MSO logic of some matroids. From a logic point of view, a matroid M on a finite ground set E is the collection of all subsets 2^E together with a unary predicate *indep*, such that *indep*(F) if and only if $F \subseteq E$ is independent in M. We denote by MS_M the language of MSO logic applied to such matroids.

See [6] for more details, and for the important result:

Theorem 3.1. (PH [6]) *Let $\mathbb{F}$ be a finite field, $t \geq 1$, and let ϕ be a sentence in the language of MS_M. There exists a finite tree automaton $\mathcal{A}_t^\phi$ such that the following is true: If an $\mathbb{F}$-represented matroid M is given together with its branch-decomposition of width at most t, then $\mathcal{A}_t^\phi$ decides whether $M \models \phi$.*

Combined with the cubic FPT algorithm [7] for approximating a branch-decomposition of an $\mathbb{F}$-represented matroid, Theorem 3.1 implies an efficient

(fixed-parameter tractable) solution of any MS_M-definable property on matroids of bounded branch-width represented over finite fields.

As we show now, the assumption of $\mathbb{F}$ being a finite field is really critical in Theorem 3.1: We describe a quite simple MS_M-definable property of a spike matroid – being a "free spike", which likely cannot be efficiently recognized on spikes represented by rational matrices (or over other infinite fields), although all spikes have trivial branch-decompositions of width 3 by Proposition 2.1 (d).

We say that a spike S on the ground set $\{e_1, \ldots, e_n, f_1, \ldots, f_n\}$ as above is a *free spike*, if every set $X \subseteq E(S)$ such that $|\{e_j, f_j\} \cap X| = 1$ for all $j \in [1, n]$ is a basis. (Simply speaking, there are no other dependencies among the elements of a free spike than those forced by the definition. There is just one free spike for each rank up to isomorphism.) Actually, when $\mathbb{F}$ is a prime finite fields, then it follows from Corollary 2.4(b) that a sufficiently large free spike cannot be represented over $\mathbb{F}$. Hence one can actually decide in constant time whether a matrix over prime $\mathbb{F}$ represents the free spike. On the other hand, we prove:

Theorem 3.2. *Let $n \geq 3$, and S be the matroid represented by a matrix $[\,\boldsymbol{I}_n \,|\, \boldsymbol{D}^1(x_1, \ldots, x_n)\,]$ over the rational numbers $\mathbb{Q}$ where $x_1, \ldots, x_n \in \mathbb{Q} - \{1\}$. Then it is NP-complete to recognize that S is* not *the rank-n free spike.*

Proof. If S is not the free spike, then by Corollary 2.4 we find out a circuit intersecting each leg of S in one element.

We reduce the problem of recognizing that S represented over $\mathbb{Q}$ is not the free spike from finding a solution to the NP-complete [3] PARTITION problem over integers. (Briefly saying, the PARTITION problem asks whether a given multiset of integers can be partitioned into two parts such that their sums equal.) Let $T = \{t_1, t_2, \ldots, t_n\}$ be a multiset of positive integers – an input to the PARTITION problem. Let $t = t_1 + t_2 + \ldots + t_n$. We denote by $z_i = \frac{-2t_i}{t}$ and $x_i = \frac{1}{z_i} + 1$, for $i \in [1, n]$. Assume $I \subset [1, n]$ is such that the multiset partition $\big(\{t_i : i \in I\}, \{t_i : i \in [1, n] - I\}\big)$ is a solution to PARTITION. That is equivalent to $\sum_{i \in I} z_i = -1$, i.e. $\sum_{i \in I} \frac{1}{x_j - 1} = -1$. Hence by Corollary 2.4 (b), S is not the free spike. The converse direction proceeds in the same way. ∎

Lemma 3.3. *There is a sentence ψ in the language of matroidal MS_M; such that ψ is true for a rank-n spike S, $n \geq 5$, if and only if S is the free spike.*

Proof. First, using Proposition 2.1 (a), we identify the legs of the spike S by the predicate $\mathit{leg}(e, f) \equiv \forall x\, \exists y \neq x \big(x = e \vee x = f \vee \mathit{circuit}(\{e, f, x, y\})\big)$, where $\mathit{circuit}(C) \equiv \neg \mathit{indep}(C) \wedge \forall D \big(D \not\subseteq C \vee D = C \vee \mathit{indep}(D)\big)$. Then S is the free spike if and only if every set intersecting each leg in at most one element is independent. Hence we write $\psi = \mathit{free\_spike} \equiv \forall F \left[\mathit{indep}(F) \vee \exists e, f \left(\mathit{leg}(e, f) \wedge e, f \in F\right)\right]$. ∎

Corollary 3.4. *Let ϕ be a sentence in the language of MS_M, and M be a matroid represented by a matrix over the rationals. Then it is NP-hard to decide whether $M \models \phi$, even if M is known to have branch-width* 3.

∎

4 On a Hard Counting Problem

The well-known class NP is the collection of all decision problems to which the answer YES can be proved in polynomial time. The class $\#P$ is the enumerative counterpart of NP; a problem of counting certain objects belongs to $\#P$ if each one of the objects is recognizable in polynomial time. (A generic example of a $\#P$-problem is counting the solutions of a given boolean formula.) We refer to Garey, Johnson [3] and to Valiant [17] for detailed formal definitions. For our purpose we need to know that the counting variant $\#$PARTITION (enumerating the number of equal-sum partitions of the given multiset of integers) is $\#P$-complete, which follows directly by the reduction used in [3].

Our interest is in the following, notoriously hard, counting invariant: The Tutte polynomial $T(G; x, y)$ of a graph G, as defined below, has many applications, not only in graph theory but also in other fields such as knot theory and statistical physics. One important feature of the Tutte polynomial is that by evaluating $T(G; x, y)$ at special points in the plane one obtains several parameters of G. For example, $T(G; 1, 1)$ is the number of spanning trees of G, and $T(G; 2, 1)$ is the number of forests of G. Jaeger, Vertigan and Welsh [11] showed that evaluating the Tutte polynomial of a graph is #P-hard at every point except those lying on the hyperbola $(x-1)(y-1) = 1$ and eight special points, including at $(1, 1)$.

Tutte's original definition of this polynomial [16] is, indeed, naturally matroidal. The *Tutte polynomial* of a matroid M on the ground set E is

$$T(M; x, y) = \sum_{A \subseteq E} (x-1)^{r_M(E)-r_M(A)} (y-1)^{|A|-r_M(A)} \; .$$

(For graphs, $T(G; x, y) = T(M(G); x, y)$.) In essence, knowing the Tutte polynomial means knowing, for each i, j, how many subsets $A \subseteq E$ are there with size i and rank j.

Although the Tutte polynomial includes many hard enumeration problems, Noble [13] has shown that its computation is fixed-parameter tractable on graphs of bounded tree-width. Our matroidal extension of that result follows.

Theorem 4.1. (PH [5]) *Let $\mathbb{F}$ be a finite field, and t an integer constant. If M is a matroid of branch-width at most t represented by a matrix over $\mathbb{F}$, then the Tutte polynomial $T(M; x, y)$ can be computed in polynomial (FPT) time.*

Next we show that there is likely no nontrivial extension of Theorem 4.1 possible in the case of an infinite field $\mathbb{F}$. (Moreover, notice that we actually show hardness of evaluating the Tutte polynomial at $(1, 1)$, which is easy for all graphs but not for matroids.)

Lemma 4.2. *Let $\mathbb{F}$ be an arbitrary field, $n \geq 3$, and $z_1, \ldots, z_n \in \mathbb{F} - \{0\}$. If b is the number of bases of the spike S represented by $[\, \boldsymbol{I}_n \,|\, \boldsymbol{D}^1(x_1, \ldots, x_n)\,]$ where $x_i = \frac{1}{z_i} + 1$, and c is the number of subsets $J \subset [1, n]$ such that $\sum_{j \in J} z_j = -1$, then*

$$b = n(n-1)2^{n-2} + 2^n - c \, .$$

Proof. By Corollary 2.4 (a), any spike S has exactly $n(n-1)2^{n-2}$ bases which contain the two elements of some leg of S – choose the leg contained in a basis, then choose another leg avoiding it, and complete with an arbitrary choice from the remaining legs. Additionally, the rank-n free spike has exactly 2^n bases which intersect each leg in one element. There are no other bases in S than described above. If S is not the free spike, then each solution to $J \subset [1,n]$: $\sum_{j\in J} z_j = -1$ reduces the number of bases by 1 by Corollary 2.4 (b). Hence the formula follows. ∎

Theorem 4.3. *Let $n \geq 5$, and S be the rank-n matroid represented by a matrix $[\boldsymbol{I}_n \mid \boldsymbol{D}^1(x_1,\ldots,x_n)]$ over the rational numbers $\mathbb{Q}$ where $x_1,\ldots,x_n \in \mathbb{Q}-\{1\}$. Then it is $\#P$-hard to compute the Tutte polynomial $T(S;x,y)$ for S.*

Proof. As in the proof of Theorem 3.2, we see that each solution to PARTITION on $T=\{t_1,t_2,\ldots,t_n\}$ is in a one-to-one correspondence with a solution to $\sum_{j\in J} z_j = -1$, where $z_i = \frac{-2t_i}{t}$. Hence by Lemma 4.2, we have a direct relation between the number of solutions to PARTITION and the number of bases of S. Thus counting all bases of S is a $\#P$-hard problem. The last step is in realizing that the evaluation $T(S;1,1)$ equals the number of bases of S. ∎

5 Extension to all Infinite Fields

We would like to extend the PARTITION problem to an arbitrary infinite field $\mathbb{F}$. If $\mathbb{F}$ has characteristic 0, then $\mathbb{F}$ contains the integers, and so the extension is trivial. So let us assume that $\mathbb{F}$ has characteristic $p>0$. That is the place where we hit a surprising obstacle — the PARTITION problem itself is trivial for $p=2$! Fortunately, the reduction for PARTITION in [3] shows an easy way around this obstacle. Another question arises about how to present the input numbers. Here we use a trick, based on the fact that infinite fields of characteristic $p>0$ may be viewed as infinite-dimensional vector spaces over $GF(p)$, with the addition and scalar multiplication as in $GF(p)$. So we consider the following problem instead.

Definition. The p-VECTORSUM problem for a prime p is defined as follows: Let $F_0 \subset GF(p)^\omega$ be the set of all ω-dimensional vectors over $GF(p)$ having finitely many nonzero coordinates.
Input: A multiset $T \subseteq F_0$ of nonzero vectors, and a nonzero vector $\boldsymbol{t} \in F_0$.
Question: Is there a subset $T_1 \subseteq T$ such that $\sum_{\boldsymbol{a}\in T_1} \boldsymbol{a} = \boldsymbol{t}$?

The size of the input multiset T is $m \cdot n$ where $|T| = n$ and m is the length of the minimal prefix covering all nonzero coordinates of vectors in T. The $\#p$-VECTORSUM problem is the natural counting counterpart of p-VECTORSUM.

Lemma 5.1. *The p-VECTORSUM problem is NP-complete for every prime p, and the corresponding $\#p$-VECTORSUM problem is $\#P$-complete.*

Proof. We will use only vectors with coordinates in $\{-1,0,1\}$. Notice that $-1=1$ if $p=2$. Let $\sum(X)$ be a shortcut for $\sum_{x\in X} x$, and let c^k mean k-times repetition of the coordinate c.

Claim 1. Let $T^o = \{(0^3,1,0^3), (0^4,1,0^2), (-1,0^3,1,0^2), (0^5,1,0), (-1,0^4,1,0), (0^6,1), (-1,0^5,1)\}$, and let $T = \{(0,1,-1,0^4), (0,1,0,-1,0^3), (0^2,1,0,-1,0^2), (0^2,1,0^2,-1,0), (0^3,1,0^2,-1), (1,0^6)\} \cup T^o$ be sets of vectors. Let $\boldsymbol{u} = (0,1,0^5)$. There are exactly 7 sets $U_1,\ldots,U_7 \subseteq T$ such that $\sum(U_i) = \boldsymbol{u}$, $i \in [1,7]$. Moreover, $U_i \cap T^o = \{\boldsymbol{u}_i\}$ for $i \in [1,7]$ where $\{\boldsymbol{u}_i : i \in [1,7]\} = T^o$.

Having Claim 1 at hand, the rest of the proof proceeds in the same way as the traditional reduction [3] of Partition to 3-SAT, except the last step. Let $\varphi = c_1 \wedge \ldots \wedge c_n$ be a SAT formula in conjunctive normal form such that each clause c_i contains exactly 3 literals, and let $x_1,\ldots,x_m$ be the variables involved in φ. We interpret satisfiability of φ via vectors from $GF(p)^{m+7n+3n}$. The coordinates of the vectors are divided into three sections; variable, clause, and incidence sections of lengths m, $7n$, and $3n$.

The incidence section of vector coordinates is indexed by the literals of all clauses of φ. For a variable x_i, $i \in [1,m]$, we construct two vectors $\boldsymbol{a}_i$, $\boldsymbol{a}'_i$, both having 1 at the position $i-1$ of the variable section. Moreover, $\boldsymbol{a}_i$ has 1 at the positions of the incidence section indexed by literals which are x_i. Similarly, $\boldsymbol{a}'_i$ has 1 at the positions of the incidence section indexed by literals which are $\neg x_i$. All remaining coordinates are set to 0. The meaning of this piece of our construction is that we choose a $True$ or $False$ value for x_i by picking $\boldsymbol{a}_i$ or $\boldsymbol{a}'_i$, respectively, in the selected subset.

The chosen valuation of $x_1,\ldots,x_m$ satisfies φ if each clause is valued $True$, i.e. if each clause triple (of coordinates) in the incidence section has a nonzero coordinate. There are $2^3 - 1 = 7$ true valuations for each clause, and that is where we apply Claim 1. For a clause c_j, $j \in [1,n]$, we make a "shifted" copy T_j of the set T from Claim 1 such that the coordinates of vectors from T now start at the position $7(j-1)$ of the clause section. The seven vectors from $T^o_j \subset T_j$, moreover, have the coordinate triples for the clause c_j in the incidence section set to $(1,1,0)$, $(1,0,1)$, $(0,1,1)$, $(1,0,0)$, $(0,1,0)$, $(0,0,1)$, $(0,0,0)$, respectively. All remaining coordinates are set to 0.

Finally, we set $\boldsymbol{v} = (1^m, \overbrace{0,1,0^6,1,\ldots,0^6,1,0^5}^{7n}, 1^{3n})$, and $U = \bigcup_{j\in[1,n]} T_j \cup \bigcup_{i\in[1,m]} \{\boldsymbol{a}_i, \boldsymbol{a}'_i\}$. Assume that $V \subseteq U$ is such that $\sum(V) = \boldsymbol{v}$. Then the variable section of the coordinates guarantees that exactly one of $\boldsymbol{a}_i, \boldsymbol{a}'_i$ for $i \in [1,m]$ belongs to V, which determines the value for x_i. The clause section guarantees that exactly one vector from each set T^o_j is selected in V for $j \in [1,n]$. And since the incidence section for the clause c_j has each coordinate 1, which cannot be obtained just using a vector from T^o_j, the clause c_j, $j \in [1,n]$ must be true in the selected valuation. The converse direction proceeds similarly. Therefore, the solutions to "$V \subseteq U$ such that $\sum(V) = \boldsymbol{v}$" are in a one-to-one correspondence with satisfying assignments to the formula φ.

The size of W is clearly polynomial in the size of φ, and we formally pad all other coordinates of the vectors with 0's. So the p-VectorSum problem for the input $U, \boldsymbol{v}$ is NP-complete. Moreover, the one-to-one correspondence be-

tween solutions to $U, \boldsymbol{v}$ and satisfying assignments to φ proves also that the #p-VectorSum problem is #P-complete. ∎

Now we are prepared to extend our results to an arbitrary infinite field $\mathbb{F}$. As above, we consider $p > 0$, and a countable vector subspace $GF(p)^\omega$ of $\mathbb{F}$. As "numbers" we mean the set $F_0 \subset GF(p)^\omega$ of all the vectors having finitely many nonzero coordinates. Notice that the zero vector in F_0 coincides with $0 \in \mathbb{F}$. We create a set F_1 of all finite-length symbolic expressions formed over F_0 using valid operations of the field $\mathbb{F}$. Formally, F_1 is the language over the alphabet $F_0 \cup \{0, 1, -, ^{-1}, +, \cdot, (,)\}$ defined recursively as follows:

- $F_0 \subset F_1$ and $0, 1 \in F_1$.
- If $a, b \in F_1$, then $-(a) \in F_1$, $(a) + (b) \in F_1$, and $(a) \cdot (b) \in F_1$.
- If $a \in F_1$ such that $\text{eval}(a) \neq 0$, then $(a)^{-1} \in F_1$.

The mapping $\text{eval} : F_1 \to \mathbb{F}$ is the homomorphism defined by $\mathbb{F}$-evaluations of the symbolic expressions. Notice that each expression in F_1 is a finite word.

The definition of F_1 may seem unnatural at the first look, but it actually is analogous to standard handling of rational numbers — algorithms computing with rational numbers usually input and work with symbolic fractions $x = \frac{a}{b}$, and not with the decimal expansion of x. Thus instead of "arithmetic in $\mathbb{F}$", we actually consider arithmetic of the symbolic expressions in F_1, and we allow symbolic input from F_1. The understanding behind this convention is that, for any particular infinite field $\mathbb{F}$, an oracle is given for computing in $\mathbb{F}$, and then the symbolic expressions my be easily parsed in time polynomial in their size. However, one should keep in mind that we have no means to decide the $\mathbb{F}$-value of a symbolic expression without such an oracle.

Theorem 5.2. *Let $n \geq 5$, and let $\mathbb{F}$ be an infinite field. If S is the matroid represented by a matrix $[\boldsymbol{I}_n \mid \boldsymbol{D}^1(x_1, \ldots, x_n)]$ over $\mathbb{F}$ where $x_1, \ldots, x_n \in \mathbb{F} - \{1\}$, then it is NP-hard to recognize that S is not the rank-n free spike.*

Proof. If $\mathbb{F}$ has characteristic 0, then $\mathbb{F}$ contains the rational numbers as a subfield, and hence the statement is proved in Theorem 3.2. Otherwise, for characteristic $p > 0$, we use a reduction from the p-VectorSum problem.

Let $F_0 \subset GF(p)^\omega$ be the set of all ω-dimensional vectors over $GF(p)$ having finitely many nonzero coordinates. Suppose that $T = \{\boldsymbol{t}_1, \boldsymbol{t}_2, \ldots, \boldsymbol{t}_n\} \subset F_0 - \{\boldsymbol{0}\}$ is a multiset of nonzero vectors, and that $\boldsymbol{t} \in F_0 - \{\boldsymbol{0}\}$. We denote by $z_i = -\boldsymbol{t}_i \cdot \boldsymbol{t}^{-1} \in F_1$ and $x_i = z_i^{-1} + 1 \in F_1$ for $i \in [1, n]$, as symbolic expressions. Then, by Corollary 2.4 (b), the matrix $[\boldsymbol{I}_n \mid \boldsymbol{D}^1(x_1, \ldots, x_n)]$ does not represent the rank-n free spike if and only if there is a subset $J \subseteq [1, n]$ such that $\sum_{j \in J} z_j = -1$ in $\mathbb{F}$, hence if $\sum_{j \in J} \boldsymbol{t}_j = \boldsymbol{t}$ in $GF(p)^\omega$.

Since we have only used symbolic operations in F_1, the reduction to a p-VectorSum instance $T, \boldsymbol{t}$ is finished in polynomial time. The reduction itself does not depend on $\mathbb{F}$ except for the characteristic p. So we are done by Lemma 5.1. ∎

In a way analogous to Theorem 5.2, we extend also Theorem 4.3 for an arbitrary infinite field $\mathbb{F}$. (We remind the reader again that the matroid S here, a spike, has branch-width three.)

Theorem 5.3. *Let $n \geq 5$, and let $\mathbb{F}$ be an infinite field. If S is the matroid represented by a matrix $[\boldsymbol{I}_n \,|\, \boldsymbol{D}^1(x_1, \ldots, x_n)]$ over $\mathbb{F}$ where $x_1, \ldots, x_n \in \mathbb{F} - \{1\}$, then it is #P-hard to compute the Tutte polynomial $T(S; x, y)$ for S.*

∎

6 Conclusions

The overall goal of our paper is to show a big difference of algoritmic behavior of matroids represented over finite and over infinite fields with respect to the structural parameter branch-width: While a large class of problems is efficiently solvable on matroids of bounded branch-width represented over finite fields, bounding the branch-width does not seem to help with solving those problems on matroids represented over infinite fields. (We remark that our bound 3 on branch-width is the smallest nontrivial value.)

The methods used here are further extended in [9] to show hardness of the problem of finding a matroid minor in a quite special setting. Moreover, inspired by a very helpful suggestion of an anonymous referee, we are applying these methods to show [10] that it is hard to decide whether a matroid (a spike) represented over $\mathbb{Q}$ has a representation over a (fixed) non-prime finite field $\mathbb{F}$.

Lastly we mention a relation of this paper to our other work [8]. In particular, [8] shows a simple interpretation of the MS_1 theory of all graphs in the MS_M theory of matroid spikes. (That is another evidence of hardness of the MS_M theory of spikes, since many classical hard graph problems like 3-colouring or dominating set are formulated in MS_1.) In general, we would like to say that branch-width is likely not a good structural parameter for all matroids, although it works nicely for matroids over finite fields. It would, of course, be nice to find a finer (than branch-width) structural parameter providing nontrivial FPT algorithms for matroids over infinite fields. See [8] for a brief outline of a possible future research in this direction.

References

1. B. Courcelle, *The Monadic Second-Order Logic of Graphs I. Recognizable sets of Finite Graphs*, Information and Computation 85 (1990), 12–75.
2. R.G. Downey, M.R. Fellows, Parametrized Complexity, Springer-Verlag New York, 1999.
3. M.R. Garey, D.S. Johnson, Computers and Intractability, W.H. Freeman and Company, New York 1979.
4. J.F. Geelen, A.H.M. Gerards, G.P. Whittle, *Branch-Width and Well-Quasi-Ordering in Matroids and Graphs*, J. Combin. Theory Ser. B 84 (2002), 270–290.
5. P. Hliněný, *The Tutte Polynomial for Matroids of Bounded Branch-Width*, Combin. Probab. Computing, to appear (accepted 2004).

6. P. Hliněný, *Branch-Width, Parse Trees, and Monadic Second-Order Logic for Matroids*, J. Combin. Theory Ser. B, to appear (accepted 2005).
7. P. Hliněný, *A Parametrized Algorithm for Matroid Branch-Width*, SIAM J. Computing 35 (2005), 259–277 (electronic).
8. P. Hliněný, D. Seese, *Trees, Grids, and MSO Decidability: from Graphs to Matroids*, Theoretical Computer Sci., to appear (accepted 2005).
9. P. Hliněný, *On Hardness of the Matroid Minor Problem*, submitted (2005).
10. P. Hliněný, *On the Complexity of the Matroid Representability Problem*, in preparation (2006).
11. F. Jaeger, D.L. Vertigan, D.J.A. Welsh, *On the Computational Complexity of the Jones and Tutte Polynomials*, Math. Proc. Camb. Phil. Soc. 108 (1990), 35–53.
12. D. Mayhew, *Matroids and Complexity*, DPhil Thesis, University of Oxford, 2005.
13. S.D. Noble, *Evaluating the Tutte Polynomial for Graphs of Bounded Tree-Width*, Combin. Probab. Computing 7 (1998), 307–321.
14. J.G. Oxley, Matroid Theory, Oxford University Press, 1992,1997.
15. J.G. Oxley, D. Vertigan, G. Whittle, *On Inequivalent Representations of Matroids over Finite Fields*, J. Combin. Theory Ser. B 67 (1996), 325–343.
16. W.T. Tutte, *A Ring in Graph Theory*, Proc. Camb. Phil. Soc. 43 (1947), 26–40.
17. L.G. Valiant, *The Complexity of Enumeration and Reliability Problems*, SIAM J. Computing 8 (1979), 410–421.