Basic Research in Computer Science

BRICS RS-96-33 Buss et al.: The Computational Complexity of Some Problems of Linear Algebra

The Computational Complexity of Some Problems of Linear Algebra

Jonathan F. Buss
Gudmund S. Frandsen
Jeffrey O. Shallit

BRICS Report Series **RS-96-33**

ISSN 0909-0878 **September 1996**

**Copyright © 1996, BRICS, Department of Computer Science
University of Aarhus. All rights reserved.**

**Reproduction of all or part of this work
is permitted for educational or research use
on condition that this copyright notice is
included in any copy.**

See back inner page for a list of recent publications in the BRICS Report Series. Copies may be obtained by contacting:

**BRICS
Department of Computer Science
University of Aarhus
Ny Munkegade, building 540
DK - 8000 Aarhus C
Denmark**

**Telephone: +45 8942 3360
Telefax: +45 8942 3255
Internet: BRICS@brics.dk**

BRICS publications are in general accessible through WWW and anonymous FTP:

`http://www.brics.dk/`
`ftp ftp.brics.dk (cd pub/BRICS)`

The Computational Complexity of Some Problems of Linear Algebra

Jonathan F. Buss* Gudmund S. Frandsen†
Jeffrey O. Shallit*

September 30, 1996

Abstract

We consider the computational complexity of some problems dealing with matrix rank. Let E, S be subsets of a commutative ring R. Let $x_1, x_2, \ldots, x_t$ be variables. Given a matrix $M = M(x_1, x_2, \ldots, x_t)$ with entries chosen from $E \cup \{x_1, x_2, \ldots, x_t\}$, we want to determine

$$\text{maxrank}_S(M) = \max_{(a_1,a_2,\ldots,a_t)\in S^t} \text{rank } M(a_1, a_2, \ldots a_t)$$

and

$$\text{minrank}_S(M) = \min_{(a_1,a_2,\ldots,a_t)\in S^t} \text{rank } M(a_1, a_2, \ldots a_t).$$

There are also variants of these problems that specify more about the structure of M, or instead of asking for the minimum or maximum

*Supported in part by grants from the Natural Sciences and Engineering Research Council (NSERC) of Canada and by the Information Technology Research Centre (ITRC) of Ontario. Address: Department of Computer Science, University of Waterloo, Waterloo, Ontario N2L 3G1, CANADA. Email: jfbuss@math.uwaterloo.ca and shallit@graceland.uwaterloo.ca

†Supported by the ESPRIT Long Term Research Programme of the EU, under project number 20244 (ALCOM-IT), and by Basic Research in Computer Science (BRICS), Centre of the Danish National Research Foundation. Address: Department of Computer Science, University of Aarhus, Ny Munkegade, DK-8000 Aarhus C, DENMARK. Email: gsfrandsen@daimi.aau.dk

rank, ask if there is some substitution of the variables that makes the matrix invertible or noninvertible.

Depending on E, S, and on which variant is studied, the complexity of these problems can range from polynomial-time solvable to random polynomial-time solvable to NP-complete to $PSPACE$-solvable to unsolvable.

1 Introduction

We consider the computational complexity of some problems of linear algebra — more specifically, problems dealing with matrix rank.

Our mathematical framework is as follows. If R is a commutative ring, then $\mathcal{M}_n(R)$ is the ring of $n \times n$ matrices with entries in R. The rows α_i of a matrix are *linearly independent* over R if $\sum_i c_i \alpha_i = 0$ (with $c_i \in R$) implies $c_i = 0$ for all i, and similarly for the columns.

The *determinant* of $M = (a_{ij})_{1 \le i,j \le n}$ is defined by

$$\det M = \sum_{P=(i_1,i_2,\ldots,i_n)} (\mathrm{sgn} P) a_{1,i_1} a_{2,i_2} \cdots a_{n,i_n},$$

where

$$P = \begin{pmatrix} 1 & 2 & \cdots & n \\ i_1 & i_2 & \cdots & i_n \end{pmatrix}$$

is a permutation of $\{1, 2, \ldots, n\}$. A matrix is invertible over R if and only if its determinant is invertible over R [10].

The *rank* of a matrix M is the maximum number of linearly independent rows. Rank can also be defined as the maximum number of linearly independent columns, and it is well-known [10] that these two definitions coincide. We denote the rank of M as rank M. An $n \times n$ matrix is invertible iff its rank is n.

A $k \times k$ *submatrix* of M is the array formed by the elements in k specified rows and columns; the determinant of such a submatrix is called a $k \times k$ *minor*. The rank of M can also be defined as the maximum size of an invertible minor.

The problems we consider are along the following lines: let E, S be two subsets of R. We are given an $n \times n$ matrix $M = M(x_1, x_2, \ldots, x_t)$ with

entries chosen from $E \cup \{x_1, x_2, \ldots, x_t\}$, where the x_i are distinct variables. We want to compute

$$\text{maxrank}_S(M) = \max_{(a_1,a_2,\ldots,a_t)\in S^t} \text{rank } M(a_1, a_2, \ldots, a_t) \quad (1)$$

$$\text{minrank}_S(M) = \min_{(a_1,a_2,\ldots,a_t)\in S^t} \text{rank } M(a_1, a_2, \ldots, a_t). \quad (2)$$

Evidently there is no need to distinguish between column rank and row rank in this definition. We do not necessarily demand that we be able to exhibit the actual t-tuple that achieves the maximum or minimum rank.

One operation that we will frequently use in this paper is taking a list of matrices $M_1, M_2, \ldots, M_k$ and constructing a large matrix M by placing each of the M_i consecutively on the main diagonal, and zeroes elsewhere. For the result we write $M = \text{diag}(M_1, M_2, \ldots, M_k)$. In this case, we have

$$\det M = \prod_{1\le i\le k} \det M_i; \quad (3)$$

$$\text{minrank}_S(M) \ge \sum_{1\le i\le k} \text{minrank}_S(M_i); \quad (4)$$

$$\text{maxrank}_S(M) \le \sum_{1\le i\le k} \text{maxrank}_S(M_i). \quad (5)$$

We will show that, depending on the arrangement of the variables in M, and on the sets E, S, the complexity of the minrank and maxrank problems ranges from being in P to being unsolvable.

There are several reasons for studying these problems. First, the problems seem — to us, at least — natural questions in linear algebra. Second, a version of the minrank problem is very closely related to determining the minimum rank rational series that approximates a given formal power series to a given order; see [7, 16] and Section 15 of the present paper. Third, the maxrank problem is related to the problem of matrix rigidity which has recently received much attention [17, 6, 11], and may help explain why good bounds on matrix rigidity are hard to obtain.

2 Some examples

Before describing our complexity results, we illustrate the minrank and maxrank problems with some examples. First, consider the matrix

$$M = \begin{bmatrix} x_1 & x_2 & 2 \\ 4 & x_1 & 4 \\ 0 & 0 & x_3 \end{bmatrix}.$$

Then $\text{minrank}(M) = 1$, attained at $(x_1, x_2, x_3) = (2, 1, 0)$. Also, $\text{maxrank}(M) = 3$, attained at $(x_1, x_2, x_3) = (2, 2, 1)$.

Note that both $\text{minrank}_S(M)$ and $\text{maxrank}_S(M)$ may depend on S. For example, if $S = \mathbb{Q}$ and

$$M = \begin{bmatrix} 1 & x \\ x & 2 \end{bmatrix}$$

then $\text{minrank}_S(M) = 2$, while if $S = \mathbb{R}$ then $\text{minrank}_S(M) = 1$, as can easily be seen by taking $x = \sqrt{2}$. Similarly if $S = \mathbb{R}$ and

$$M = \begin{bmatrix} 1 & x \\ x & -2 \end{bmatrix}$$

then $\text{minrank}_S(M) = 2$, while if $S = \mathbb{C}$ then $\text{minrank}_S(M) = 1$, as can be seen by taking $x = \sqrt{2}i$. Clearly $\text{minrank}_S(M) \geq \text{minrank}_{S'}(M)$ if $S \subseteq S'$.

A similar phenomenon occurs for the maxrank problem. For example, if $S = GF(2)$, the finite field with two elements, and

$$M = \begin{bmatrix} x & x \\ 1 & x \end{bmatrix},$$

then $\text{maxrank}_S(M) = 1$. On the other hand, if $S = GF(4)$, then $\text{maxrank}_S(M) = 2$, as can be seen by taking $x = \alpha$, where α is a generator of the multiplicative group of S. Clearly $\text{maxrank}_S(M) \leq \text{maxrank}_{S'}(M)$ if $S \subseteq S'$.

3 Summary of Results

Most of our complexity results for the computation of minrank and maxrank are naturally phrased in terms of the decision problems given in Table 1.

Fixed: R, a commutative ring.
$E, S \subseteq R$.
Input: M, an $n \times n$ matrix with entries from $E \cup \{x_1, \ldots, x_t\}$.
k, a non-negative integer.

Problem	Input	Decide
MINRANK	M, k	$\min_{(a_1,\ldots,a_t)\in S^t}$ rank $M(a_1, \ldots a_t) \leq k$?
MAXRANK	M, k	$\max_{(a_1,\ldots,a_t)\in S^t}$ rank $M(a_1, \ldots a_t) \geq k$?
SING	M	$\exists (a_1, \ldots, a_t) \in S^t$ such that $\det M(a_1, \ldots a_t) = 0$?
NONSING	M	$\exists (a_1, \ldots, a_t) \in S^t$ such that $\det M(a_1, \ldots a_t) \neq 0$?

Table 1: Decision problems.

S	E	MAXRANK NONSING	SING	MINRANK
$GF(q)$	$\{0,1\} \subseteq E \subseteq GF(q)$	*NP*-complete		
$\mathbb{Z}$	$\{0,1\} \subseteq E \subseteq \mathbb{Q}$	*RP*	r.e.; undecidable	
$\mathbb{Q}$			r.e.; *NP*-hard	
$\mathbb{R}$			*PSPACE*; *NP*-hard	
$\mathbb{C}$				

Table 2: Complexity bounds for decision problems.

We have introduced two special problems, SING(ularity) and NONSING(ularity), which could possibly be easier than the more general min-rank/maxrank problems.

Table 2 summarizes our results on the complexity of the four decision problems. We put the problems MAXRANK and NONSING together, since we have not been able to separate their complexities, although we do not know whether they have the same complexity in general. We have good evidence that the MINRANK and SING problems do not in general have the same complexity. Over $\mathbb{C}$, the MINRANK problem is *NP*-hard (Section 11), whereas SING has a random polynomial-time solution (Section 5).

The exact value of E is not important for our bounds. All our lower bounds are valid for $E = \{0, 1\}$ and all our upper bounds are valid when

E is $\mathbb{Q}$ or a finite-dimensional field extension of $\mathbb{Q}$ (respectively, when E is $GF(q)$ or a finite-dimensional field extension of $GF(q)$, when the characteristic is finite). For the upper bounds, we assume the input size to be the total number of bits needed to list each separate entry of the matrix M, representing numbers using the standard binary representation, representing constants in a finite-dimensional algebraic extension by arithmetic modulo an irreducible polynomial, and representing polynomials by the vectors of their coefficients. The upper bounds are also robust in another sense. We can allow entire multivariate polynomials (with coefficients from E) in a single entry of the matrix M and still preserve our upper bounds, provided such a multivariate polynomial is specified by an arithmetic formula using binary multiplication and binary addition, but no power symbol, so that the representation length of a multivariate polynomial is at least as large as its degree.

S is significant for the complexity, as shown in Table 2. However, our upper and lower bounds for $S = C$ are valid for S being any algebraically closed field (if S has finite characteristic, E must also, of course).

The results of Table 2 fall in three groups according to the proof technique used. The random polynomial-time upper bounds use a result due to Schwartz [15]. The undecidability result for $\mathbb{Z}$ uses a combination of Valiant's result that the determinant is universal [18] and Matiyasevich's proof that Hilbert's Tenth Problem is unsolvable [12]. All the remaining problems of the result table (those that are not marked either *RP* or *undecidable*) are equivalent (under polynomial-time transformations) to deciding the existential first-order theory over the field S. The equivalence implies the *NP*-hardness of all these problems, and lets us use results by Ierardi [9] and Canny [3] to obtain the *PSPACE* upper bounds for $\mathbb{C}$ and $\mathbb{R}$, respectively. Since it is presently an open problem whether the existential first-order theory over $\mathbb{Q}$ is decidable or not, we suspect it will be difficult to determine the decidability status of `MINRANK` and `SING` over $\mathbb{Q}$.

We also consider the special case when each variable in the matrix occurs exactly once. None of our lower bound proofs are valid under this restriction, and we have improved some of the upper bounds. See Table 3 for a summary. The improved upper bounds all rely on the determinant polynomial being multi-affine when no variable occurs twice. In such a case the *RP*-algorithm for singularity over $\mathbb{C}$ can be generalized to work for singularity over any

<table>
<tr><th>S</th><th>E</th><th>MAXRANK
NONSING</th><th>SING</th><th>MINRANK</th></tr>
<tr><td>$GF(q)$</td><td>$GF(q)$</td><td colspan="2" rowspan="2"></td><td>NP</td></tr>
<tr><td>$\mathbb{Z}$</td><td rowspan="4">$\mathbb{Q}$</td><td colspan="2">r.e.</td></tr>
<tr><td>$\mathbb{Q}$</td><td colspan="2" rowspan="2">RP</td><td rowspan="2">$PSPACE$</td></tr>
<tr><td>$\mathbb{R}$</td></tr>
<tr><td>$\mathbb{C}$</td><td></td><td></td></tr>
</table>

Table 3: Upper bounds for decision problems, when each variable occurs exactly once.

field.

For a very special kind of matrix, viz. row-partitionable matrices where each variable occurs exactly once, we give in Section 15 a polynomial time algorithm for computing the minimum possible rank. The algorithm works in the case where S is any field.

Since minrank is at least *NP*-hard to compute over $\mathbb{Z}$ or a field, one might consider the existence of an efficient approximation algorithm. Suppose, however, that for some fixed S (S being $\mathbb{Z}$ or a field) and $E = \{0, 1\}$, there is a polynomial time algorithm that when given matrix $M = M(x_1, \ldots, x_t)$ always returns a vector $(a_1, \ldots, a_t) \in S^t$ satisfying rank $(M(a_1, \ldots, a_t)) \leq (1 + \varepsilon) \cdot \text{minrank}_S(M)$. Then the assumption $P \neq NP$ implies $\varepsilon \geq \frac{7}{1755} \approx 0.0039886$, as we prove in Section 13. The proof uses reduction from `MAXEXACT3SAT`; i.e., we use a known nonapproximability result for `MAXEXACT3SAT` [1] combined with a *MAXSNP*-hardness proof for the minrank approximation problem.

4 Computing maxrank over infinite fields

In this section we show how to compute maxrank with a (Monte-Carlo) random polynomial-time algorithm over any infinite field.

We will also show that to solve the problem for $R = S = F$, it suffices to consider the case $R = S = \mathbb{Z}$, when F contains $\mathbb{Z}$.

Our main tool is the following lemma, adapted from a paper of Schwartz [15]:

Lemma 1 *Let $p(x_1, x_2, \ldots, x_t)$ be a multivariate polynomial of total degree at most d which is not the zero polynomial, and let F be a field containing at least $2d$ distinct elements. Then if V is any set of $2d$ distinct elements of F, $p(a_1, a_2, \ldots, a_t) = p(\mathbf{a}) \neq 0$ for at least 50% of all $\mathbf{a} \in V^d$.*

Theorem 2 *Let $M = M(x_1, x_2, \ldots, x_t)$ be a $n \times n$ matrix with entries in $F \cup \{x_1, x_2, \ldots, x_t\}$. Let $V \subseteq F$ be a set of at least $2n$ distinct elements (If $\mathbb{Z} \subseteq F$ then $V = \{-n, 1-n, \ldots, -1, 0, 1, 2, \ldots, n\}$ may be used). Choose a t-tuple $(a_1, a_2, \ldots, a_t) \in V^t$ at random. Then with probability at least $1/2$, we have*

$$\text{maxrank}_F(M) = \text{rank } M(a_1, a_2, \ldots, a_t).$$

Proof. Suppose $\text{maxrank}_F(M) = k$. Then there exists some t-tuple $(a_1, a_2, \ldots, a_t) \in F^t$ such that rank $M(a_1, a_2, \ldots, a_t) = k$. Hence, in particular, there must be some $k \times k$ minor of $M(a_1, a_2, \ldots, a_t)$ with nonzero determinant. Consider the corresponding $k \times k$ submatrix M' of $M(x_1, x_2, \ldots, x_t)$. Then the determinant of M', considered as a multivariate polynomial p in the indeterminates $x_1, x_2, \ldots, x_t$, cannot be identically zero (since it is nonzero when $x_1 = a_1, \ldots, x_t = a_t$). It now follows from Lemma 1 that p is nonzero for at least half of all elements of V^t. Thus for at least half of all these t-tuples $(a_1, a_2, \ldots, a_t)$, the corresponding $k \times k$ minor of M must be nonzero, and hence $M(a_1, a_2, \ldots, a_t)$ has rank at least k. Since $\text{maxrank}_F(M) = k$, it follows that rank $M(a_1, a_2, \ldots, a_t) = k$ for at least half of the choices $(a_1, a_2, \ldots, a_t) \in V^t$. ■

The theorem implies a random polynomial-time algorithm to compute $\text{maxrank}_F(M)$ over an infinite field F. Choose r t-tuples of the form $(a_1, a_2, \ldots, a_t)$ independently at random, and compute rank $M(a_1, a_2, \ldots, a_t)$ for each of them, obtaining ranks $b_1, b_2, \ldots, b_r$. Then with probability at least $1 - 2^{-r}$, we have $\text{maxrank}_F(M) = \max_{1 \leq i \leq r} b_i$.

It also follows from Theorem 2 that over an infinite field F, the quantity $\text{maxrank}(M)$ cannot change when we consider an extension field F' with $F \subseteq F'$, or when we consider an infinite subset $S \subseteq F$. The algorithm runs

exactly the same way so long as $V \subseteq S \subseteq F \subseteq F'$. In particular, if F has characteristic zero, $\text{maxrank}_F(M) = \text{maxrank}_{\mathbb{Z}}(M)$. Therefore the theorem also implies that the decision problem MAXRANK is in the complexity class RP for $E = \mathbb{Q}$ and $\mathbb{Z} \subseteq S$.

5 The singularity problem over an algebraically closed field

In this section we consider the complexity of the decision problem SING in the case $R = S = F$, where F is an algebraically closed field. We will show that in this case, SING $\in RP$. First, we prove the following lemmas.

Lemma 3 *Let $p(x_1, x_2, \ldots, x_t)$ be a multivariate polynomial over an infinite field F. Then p is identically zero iff p is the zero polynomial.*

Proof. If p is the zero polynomial, then the result is clear.

Otherwise assume p is not the zero polynomial. We prove the result by induction on t, the number of variables. If $t = 1$, then p is a univariate polynomial of degree d for some $d \geq 1$. This polynomial has at most d zeroes, and since F is infinite, $p(a) \neq 0$ for all but finitely many $a \in F$.

Now assume the result is true for all $t < k$; we prove it for $t = k$. Choose a variable x in p that occurs with highest degree, say d, and write p as a polynomial in x with multivariate coefficients, say $p = z_d x^d + \cdots + z_1 x + z_0$. Since p is nonconstant, we have $d \geq 1$. Now z_d is a polynomial in $k - 1$ in variables that is not the zero polynomial; hence by induction z_d is not identically zero. Choose an assignment to the variables such that $z_d \neq 0$, and call the new polynomial $q = q(x)$. Then q is not the zero polynomial, and hence by induction is not identically zero. ■

Lemma 4 *Let $p(x_1, x_2, \ldots, x_t)$ be a nonconstant multivariate polynomial over a field F. Then if F is algebraically closed, p takes on all values in F.*

Proof. We prove the result by induction on t, the number of variables. If $t = 1$, then $p = p(x)$ is a nonconstant univariate polynomial. To show p takes on all values in F, consider the equation $p(x) - c = 0$ for $c \in F$. Since F is

algebraically closed, this equation has a solution $x = x_0$, and then $p(x_0) = c$. Since c was arbitrary, the result follows.

Now consider the case $t > 1$. Write $p = y_1 + y_2 + \cdots + y_r$, where each y_i is a (possibly constant) monomial of the form $a_i x_1^{e_{i1}} x_2^{e_{i2}} \cdots x_t^{e_{it}}$. Furthermore, assume that all terms are collected, so that we never have

$$i \neq j \text{ and } (e_{i1}, e_{i2}, \ldots, e_{it}) = (e_{j1}, e_{j2}, \ldots, e_{jt}). \tag{6}$$

Choose a term y_i in which some variable, say x, occurs in the form x^e, and e is as large as any exponent occurring in any monomial of p. Since p is nonconstant, we must have $e \geq 1$. Now think of p as a polynomial in x with multivariate coefficients, and write $p = z_e x^e + \cdots + z_1 x + z_0$, where each z_i is a polynomial in the remaining variables. We claim that z_e is not the zero polynomial; if it were, then (6) would be violated. Hence by Lemma 3 there is some assignment to the variables in z_e that makes it nonzero. Make this assignment to all variables in p; the result is a nonconstant polynomial in x, and the argument for $t = 1$ then applies. ■

Theorem 5 *If $R = S = F$, and F is algebraically closed, then* SING $\in RP$.

Proof. Consider the following algorithm: Let $V \subseteq F$ be a set of at least $2n$ distinct elements (if $\mathbb{Z} \subseteq F$ then $V = \{-n, 1-n, \ldots, -1, 0, 1, 2, \ldots, n\}$ may be used). Choose r t-tuples $\mathbf{a}_1$, $\mathbf{a}_2$, ..., $\mathbf{a}_r$ at random from V^t, and evaluate the determinant $\det M(\mathbf{a}_i)$ for $1 \leq i \leq r$. If at least two different values are obtained, return "yes". If all the values obtained are the same, and all are nonzero, return "no". If all the values are the same, and all are zero, return "yes".

We claim that if there exists a t-tuple $\mathbf{a}$ such that $\det M(\mathbf{a}) = 0$, then this algorithm returns the correct result with probability at least $1 - 1/2^{r-1}$, while if there is no such t-tuple, the algorithm always returns the correct result.

To prove the claim, define $p(x_1, x_2, \ldots, x_t) = \det M(x_1, x_2, \ldots, x_t)$, a multivariate polynomial. If p is nonconstant, then by Lemma 4 it takes on all values in F, including 0. If p is constant and nonzero, then it cannot take on the value 0. Finally, if p is constant and zero, then it clearly takes on the value 0.

It now follows that our algorithm always returns the correct result except possibly when all the values obtained are the same and nonzero. In this

case we return "no", whereas if we are unlucky the answer could possibly be "yes". However, if the polynomial p is not the constant polynomial, then the polynomial $p - p(\mathbf{a}_1)$ is nonzero, and by Lemma 1 we know $p(\mathbf{a}_i) \neq p(\mathbf{a}_1)$ with probability at least $1/2$ for $2 \leq i \leq r$. It follows that the probability of making an error in this case is bounded by $1/2^{r-1}$. ■

6 Universality of the determinant

In this section, we prove a result that underlies all our lower bounds for the singularity and minrank problems: that any multivariate polynomial is the determinant of a fairly small matrix. The result was first proven by Valiant [18], but since we need a slightly modified construction and the result is fundamental to our lower bound proofs, we make this paper self-contained and give the details of the construction.

To state the result, we need a few definitions. Let an *arithmetic formula* F be a well-formed formula using constants, variables, the unary operator $\{-\}$ and the binary operators $\{+, \cdot\}$. The *length* of a formula F (denoted by $|F|$) is defined as the total number of occurrences of constants, variables and operators. For example

$$|3xy - z - 3| = |3 \cdot x \cdot y + (-(z)) + (-(3))| = 11$$

and

$$|3(x + y - 4) + 5z| = |3 \cdot (x + y + (-(4))) + 5 \cdot z| = 12.$$

(Note that our definition of formula length is not the same as Valiant's.)

Proposition 6 *Let R be a commutative ring. Let F be an arithmetic formula using constants from $E \subseteq R$ and variables from $\{x_1, \ldots, x_t\}$.*

For some $n \leq |F| + 2$, we may in time $n^{O(1)}$ construct an $n \times n$ matrix M with entries from $E \cup \{0, 1\} \cup \{x_1, \ldots, x_t\}$ such that $p_F = \det M$ and $\mathrm{minrank}_R(M) \geq n - 1$, where p_F denotes the polynomial described by formula F.

Proof. We use a modified version of Valiant's construction [18]. The main difference is that we insist that the rank of the constructed $n \times n$ matrix cannot be less than $n - 1$ under any substitution for the variables. We

also consider the negation operation explicitly, which allows us to avoid the use of negative constants in the formula, when wanted. Our construction is essentially a modification of Valiant's construction to take care of these extra requirements combined with a simplification that leads to matrices of somewhat larger size than Valiant's original construction.

Let a formula F be given. The construction falls in two parts. In the first part, we construct a series-parallel s-t-graph G_F with edge weights from $E \cup \{1\} \cup \{x_1, \ldots, x_t\}$ by induction on the structure of F as sketched in Figure 1. To such a series-parallel s-t-graph G_F, we associate the polynomial

$$p(G_F) = \sum_{\pi \text{ is } s\text{-}t\text{-path in } G_F} (-1)^{\text{length}(\pi)} \cdot \prod_{e \text{ an edge of } \pi} \text{weight}(e).$$

By induction in the structure of F, one may verify that $p_F = p(G_F)$.

In the second part of the construction, we change G_F into a cyclic graph G'_F by adding an edge from t to s of weight 1 and adding self-loops with weight 1 to all vertices different from s. The matrix $M = \{m_{ij}\}$ is simply the weight matrix for G'_F; i.e., m_{ij} is the weight of the edge from vertex i to vertex j if it exists and $m_{ij} = 0$ otherwise. The determinant of M is a sum of monomials, where each monomial is the product of the weights in a specific cycle cover of G'_F (with sign ± 1 depending on the length of the cycles). But because of the special form of G'_F each cycle cover will consist of a number of self-loops (possibly zero) and a single cycle arising from an s-t-path in G_F combined with the added edge from t to s. Hence, each s-t-path in G_F gives rise to one monomial in $\det M$, and the sign of the monomial will be -1 if and only if the path has odd length. Thus $\det M = p(G_F) = p_F$.

To see the lower bound on minrank, consider the $(n-1) \times (n-1)$ submatrix M' of M arising from erasing the column and row corresponding to the vertex s. The determinant of M' has one monomial for each cycle cover of $G'_F - \{s\}$. However, removing the vertex s breaks all cycles corresponding to paths from s to t in G_F, but with s removed all the remaining vertices have a self loop, so there is precisely one cycle cover and it consists of all the self-loops. Since all the self-loops have weight 1, we find that $\det M' = 1$, so $\text{minrank}_R(M) \geq n - 1$.

The bound $2 + |p_F|$ on the size of G_F arises because the graph G_F has in addition to the vertices s and t at most one vertex for each application of a rewrite rule from Figure 1.

Formula F	The series-parallel s-t-graph G_F with edge weights
Constant c	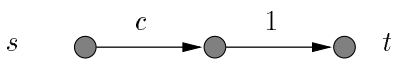
Variable x	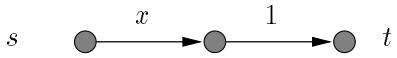
$F = -F_1$	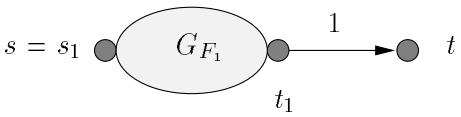
$F = F_1 \cdot F_2$	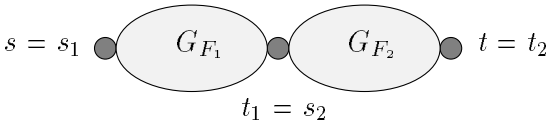
$F = F_1 + F_2$	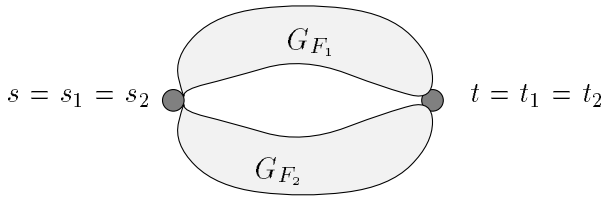

Figure 1: Inductive construction of G_F.

$$F \ : \ x_1(x_2 - 4x_3 + x_4) + x_5$$

G_F :

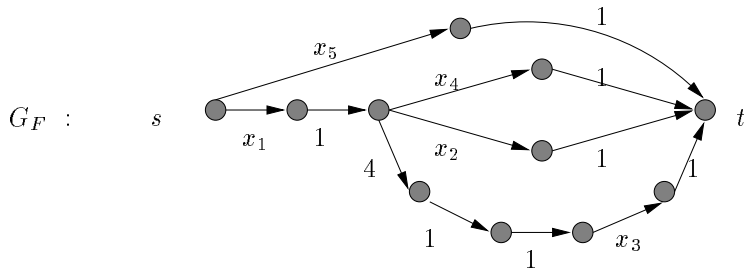

G'_F :

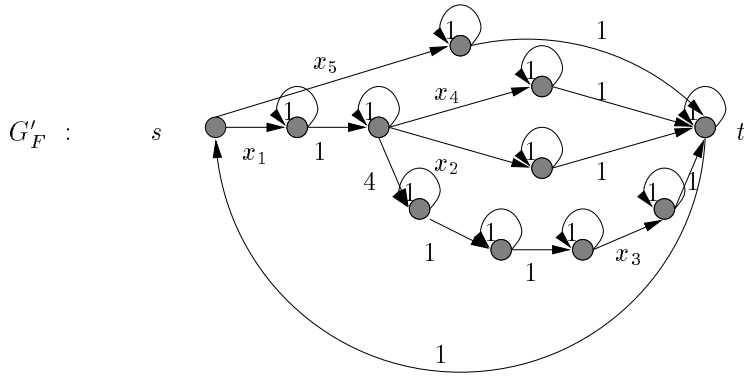

$$M = \begin{bmatrix} 0 & x_1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & x_5 & 0 \\ 0 & 1 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & x_4 & x_2 & 4 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & x_3 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 0 & 1 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix}$$

$$\det M = x_1(x_2 - 4x_3 + x_4) + x_5$$

Figure 2: Constructing a matrix with specified determinant.

Figure 2 illustrates the construction given in this proof on a specific example. ■

7 The singularity problem over the integers

In this section we prove that the decision problem SING is unsolvable for $S = \mathbb{Z}$ and $E = \{0, 1\}$.

Theorem 7 *(Undecidability of* SING *over* $\mathbb{Z}$*)*
Given a matrix $M = M(x_1, \ldots, x_t)$ *with entries from* $\{0, 1\} \cup \{x_1, \ldots, x_t\}$, *it is undecidable whether there exist* $a_1, \ldots, a_t \in \mathbb{Z}$ *such that* $\det M(a_1, \ldots, a_t) = 0$

Proof. We reduce from Hilbert's Tenth Problem [12]. An instance of Hilbert's Tenth Problem is a Diophantine equation $p(x_1, \ldots, x_t) = 0$, where p is a multivariate polynomial with integer coefficients. We construct a formula for p using only $+, -, \cdot, 0, 1$ in addition to the indeterminates by replacing each integer constant $c \geq 2$ having binary representation $c = \sum_{i=0}^{l} b_i 2^i$ with the formula

$$b_0 + (1+1)[b_1 + (1+1)[b_2 + (1+1)[b_3 + \cdots + (1+1)[b_l] \cdots]]].$$

By the construction of Proposition 6, the resulting formula f_p for the polynomial $p(x_1, \ldots, x_t)$ is turned into a matrix $M = M(x_1, \ldots, x_t)$ such that $\det M(x_1, \ldots, x_t) = p(x_1, \ldots, x_t)$. The assertion of the theorem follows from the undecidability of Hilbert's Tenth Problem. ■

8 Existential first-order theories

In this section, we describe the syntax of existential first order theories over fields and state some complexity results for the corresponding decision problems. We will apply this later to our rank problems.

For any field F, we have arithmetic operations $+, \cdot$, constants $0, 1$ and equality relation $=$. Adding the Boolean operations $\wedge, \vee, \neg$ and the existential quantifier $\exists$, we get the first order language specified by the following grammar. (Note that we require all quantifiers to be collected in a prefix to

the formula, thereby avoiding implicit universal quantification and alternation of quantifiers.)

$$\begin{array}{lcl}
\mathtt{V} & ::= & x_1 \mid x_2 \mid x_3 \mid \cdots \mid x_n \mid \cdots \\
\mathtt{C} & ::= & 0 \mid 1 \\
\mathtt{AT} & ::= & \mathtt{V} \mid \mathtt{C} \\
\mathtt{T} & ::= & \mathtt{AT} \mid (\mathtt{T} + \mathtt{T}) \mid (\mathtt{T} \cdot \mathtt{T}) \\
\mathtt{AF} & ::= & \mathtt{T} = \mathtt{T} \\
\mathtt{BF} & ::= & \mathtt{AF} \mid (\neg \mathtt{BF}) \mid (\mathtt{BF} \wedge \mathtt{BF}) \mid (\mathtt{BF} \vee \mathtt{BF}) \\
\mathtt{F} & ::= & \mathtt{BF} \mid \exists\ \mathtt{V}\ \mathtt{F}
\end{array}$$

A *sentence* is a formula with no free variables (all variables are bound by quantifiers).

We say that sentence φ is true in the field F (the field F is a *model* of the sentence φ), if the sentence evaluates to true, when quantifications are interpreted over elements in F, and arithmetic operations and constants are given the natural interpretations, and we write

$$F \models \varphi.$$

For a more formal definition of the semantics, see, for example, Enderton [5].

Examples:

$$\begin{array}{rcl}
GF(2) & \not\models & \exists x.\ x^2 + x + 1 = 0; \\
\mathbb{Q} & \models & \exists x \exists y.\ xy = 1 \ \wedge\ x^2 \neq 1; \\
\mathbb{R} & \not\models & \exists x \exists y.\ (1-x)y = 1 \ \wedge\ x(1-y) = 1; \\
\mathbb{C} & \models & \exists z.\ z^2 + 1 = 0.
\end{array}$$

The examples use both squaring and subtraction, which are shorthands for more complicated formulas using $\{+, \cdot\}$ only. For example,

$$\exists x \exists y.\ (1-x)y = 1 \ \wedge\ x(1-y) = 1$$

is shorthand for

$$\exists x \exists y \exists x' \exists y'.\ (1+x')y = 1 \ \wedge\ x(1+y') = 1 \ \wedge\ x + x' = 0 \ \wedge\ y + y' = 0.$$

For a field F, we define the existential theory of F:

$$ETh(F) = \{\varphi \ :\ F \models \varphi\}.$$

The decision problem for $ETh(F)$ is: on input φ, decide whether $F \models \varphi$.

F	Upper bound on $ETh(F)$	reference
$GF(q)$	NP	
$\mathbb{Q}$	recursively enumerable	
$\mathbb{Q}_p$	doubly exponential space	Egidi, 1993 [4]
$\mathbb{R}$	$PSPACE$	Canny, 1988 [3]; Renegar, 1992 [14]
$\mathbb{C}$	$PSPACE$	Ierardi, 1989 [9]

Table 4: Upper bounds on deciding $ETh(F)$

Proposition 8 *For F being any fixed field, $ETh(F)$ is NP-hard.*

Proof. We reduce from `3SAT`. Let C be an instance of `3SAT`; i.e.,

$$C \equiv C_1 \wedge C_2 \wedge \cdots \wedge C_k$$

where $C_i \equiv (l_{i1} \vee l_{i2} \vee l_{i3})$ and $l_{ij} \in \{y_1, y_2, \ldots, y_t\} \cup \{\overline{y_1}, \overline{y_2}, \ldots, \overline{y_t}\}$. We modify C to be an arithmetic formula f_C by replacing each y_i with the atomic formula $x_i = 1$ and replacing each $\overline{y_i}$ with the atomic formula $x_i = 0$. Clearly,

$$C \text{ is satisfiable iff } F \models \exists x_1\, \exists x_2 \cdots \exists x_t \,.\, f_C.$$

The NP-hardness follows from the NP-hardness of `3SAT`. ■

The complexity of deciding $ETh(F)$ seems to depend on the field F. Table 4 summarizes the upper bounds that we are aware of.

$ETh(GF(q))$ is in NP for any fixed finite field ($GF(q)$), since one may replace the variables with nondeterministically chosen field elements and evaluate the resulting variable free formula in polynomial time.

Similarly, $ETh(\mathbb{Q})$ is recursively enumerable, but to the best of our knowledge it is still an open problem whether $ETh(\mathbb{Q})$ is in fact decidable.

The doubly exponential space bound for the field of p-adic numbers, $\mathbb{Q}_p$ (for some fixed prime p) is proven for a more general theory than the one considered here. It is quite conceivable that a better bound can be found for our existential sentences.

One may get a $PSPACE$ bound for $\mathbb{C}$ as a corollary to the $PSPACE$ bound for $\mathbb{R}$, since arithmetic in $\mathbb{C}$ can be represented by arithmetic on pairs

$F = GF(q)$	Rewrite rules
Step 1	$t(\mathbf{x}) = 0 \quad \rightarrow \quad t(\mathbf{x})^{q-1} = 0$
Step 2	$\neg t(\mathbf{x}) = 0 \quad \rightarrow \quad 1 - t(\mathbf{x}) = 0$
	$(t_1(\mathbf{x}) = 0) \vee (t_2(\mathbf{x}) = 0) \quad \rightarrow \quad t_1(\mathbf{x}) \cdot t_2(\mathbf{x}) = 0$
	$(t_1(\mathbf{x}) = 0) \wedge (t_2(\mathbf{x}) = 0) \quad \rightarrow \quad 1 - (1 - t_1(\mathbf{x})) \cdot (1 - t_2(\mathbf{x})) = 0$
Step 3	$t(\mathbf{x}) = 0 \quad \rightarrow \quad \det M'(\mathbf{x}) = 0$

Table 5: Transforming an existential sentence to a singularity problem.

of numbers in $\mathbb{R}$. However, the proof of Ierardi [9] uses a different technique and holds for any algebraically closed field.

9 Decision problems over finite fields

In this section, we prove that both the singularity problem and the nonsingularity problem over a fixed finite field are as hard as deciding the corresponding existential first-order theory. In particular, all four decision problems that we defined are *NP*-hard (and *NP*-complete).

Lemma 9 *Let $F = GF(q)$ be a fixed finite field.*

Given an existential sentence $\exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t)$ of length m, we can in time $n^{O(1)}$ construct two $n \times n$ matrices M' and M'' with entries from $\{0, 1\} \cup \{x_1, \ldots, x_t\}$, where $n = O(mq)$ such that

$$\exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t) \quad \textit{iff} \quad \exists (a_1, \ldots, a_t) \in F^t.\ \det M'(a_1, \ldots, a_t) = 0$$

and

$$\exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t) \quad \textit{iff} \quad \exists (a_1, \ldots, a_t) \in F^t.\ \det M''(a_1, \ldots, a_t) \neq 0.$$

Proof. To construct matrix M', we modify the unquantified formula φ using the rewriting rules of Table 5.

Initially, we may assume that each atomic logic formula is on the form $t(\mathbf{x}) = 0$, for some arithmetic term $t(\mathbf{x})$. In step 1, we use the fact that over

$F = GF(q)$	Rewrite rules
Step 1	$t(\mathbf{x}) = 0 \quad \rightarrow \quad 1 - t(\mathbf{x})^{q-1} \neq 0$
Step 2	$\neg t(\mathbf{x}) \neq 0 \quad \rightarrow \quad 1 - t(\mathbf{x}) \neq 0$
	$(t_1(\mathbf{x}) \neq 0) \vee (t_2(\mathbf{x}) \neq 0) \quad \rightarrow \quad 1 - (1 - t_1(\mathbf{x})) \cdot (1 - t_2(\mathbf{x})) \neq 0$
	$(t_1(\mathbf{x}) \neq 0) \wedge (t_2(\mathbf{x}) \neq 0) \quad \rightarrow \quad t_1(\mathbf{x}) \cdot t_2(\mathbf{x}) \neq 0$
Step 3	$t(\mathbf{x}) \neq 0 \quad \rightarrow \quad \det M''(\mathbf{x}) \neq 0$

Table 6: Transforming an existential sentence to a nonsingularity problem

the field $GF(q)$, the function $x \mapsto x^{q-1}$ maps 0 to 0 and maps any nonzero number to 1.

Following step 1, we may assume that any arithmetic term takes only values in $\{0, 1\}$ under all possible assignments to variables. This assumption should make the correctness of the three rewrite rules in step 2 obvious.

When no more rewrite rules from step 2 are applicable, we have compressed $\varphi(\mathbf{x})$ to an equivalent atomic formula $t(\mathbf{x}) = 0$. In step 3, we construct a matrix M' such that $\det M' = t(\mathbf{x})$ using Proposition 6.

When using the rewriting rules, any arithmetic term occurring on the right hand side of a rule is an arithmetic formula and should stay a formula; i.e., it should not be expanded into a sum of monomials, since such a sum could be exponentially large.

The construction of matrix M'' is completely analogous, using the rewrite rules of Table 6. ■

Corollary 10 *Let F be a fixed finite field $GF(q)$. For $S = F$ and $\{0, 1\} \subseteq E \subseteq GF(q)$, the decision problems* `MAXRANK`*,* `NONSING`*,* `MINRANK` *and* `SING` *are all NP-complete.*

Proof. Clearly, these problems are in *NP*, since we may nondeterministically guess an assignment to the variables, and compute the rank of the resulting constant matrix in polynomial time.

The *NP*-hardness follows from Lemma 9 combined with Proposition 8. ■

$F = \mathbb{Q}$ or $\mathbb{R}$	Rewrite rules
Step 1	$\neg(F_1 \wedge F_2) \rightarrow (\neg F_1) \vee (\neg F_2)$
	$\neg(F_1 \vee F_2) \rightarrow (\neg F_1) \wedge (\neg F_2)$
Step 2	$\neg t(\mathbf{x}) = 0 \rightarrow 1 - z \cdot t(\mathbf{x}) = 0$
Step 3	$t(\mathbf{x}') = 0 \rightarrow t(\mathbf{x}')^2 = 0$
Step 4	$(t_1(\mathbf{x}') = 0) \vee (t_2(\mathbf{x}') = 0) \rightarrow t_1(\mathbf{x}') \cdot t_2(\mathbf{x}') = 0$
	$(t_1(\mathbf{x}') = 0) \wedge (t_2(\mathbf{x}') = 0) \rightarrow t_1(\mathbf{x}') + t_2(\mathbf{x}') = 0$
Step 5	$t(\mathbf{x}') = 0 \rightarrow \det M(\mathbf{x}') = 0$

Table 7: Transforming an existential sentence to a singularity problem, over $\mathbb{Q}$ and $\mathbb{R}$.

10 Lower bounds for singularity over $\mathbb{Q}$ and $\mathbb{R}$.

In this section, we prove that the singularity problem over either of the fields $\mathbb{Q}$ and $\mathbb{R}$ is as hard as deciding the corresponding existential first-order theory. In particular, the problems are *NP*-hard.

Lemma 11 *Let F be either of the fields $\mathbb{Q}$ or $\mathbb{R}$.*

Given an existential sentence $\exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t)$ of length m, we can in time $n^{O(1)}$ construct an $n \times n$ matrix M with entries from $\{0, 1\} \cup \{x_1, \ldots, x_{t'}\}$, where $n = O(m)$ such that

$$\exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t) \quad \textit{iff} \quad \exists (a_1, \ldots, a_{t'}) \in F^{t'}.\ \det M(a_1, \ldots, a_{t'}) = 0.$$

Proof. The proof is analogous to the proof of Lemma 9, but we handle negation differently.

To construct the matrix M, we modify the unquantified formula φ using the rewriting rules of Table 7.

Steps 3-5 in Table 7 correspond closely to steps 1-3 in Table 5, except that we have no rule for negation. The first two steps of Table 7 serve to remove negation.

In step 1, we use de Morgan's laws to move all negations down so that they are applied directly to the atomic formulas.

In step 2, we replace each negated atomic formula by an unnegated formula. We introduce a new variable z for each such atomic formula, which

represents the inverse of the term $t(\mathbf{x})$. These new variables must be existentially quantified.

In step 3, we use the fact that over each of the fields $\mathbb{Q}$ and $\mathbb{R}$, the function $x \mapsto x^2$ maps 0 to 0 and maps any nonzero number to a positive number.

Following step 3, we may assume that any arithmetic term takes only nonnegative values under all possible assignments to the variables. This assumption should make the correctness of the two rewrite rules in step 4 obvious.

When no more rewrite rules from step 4 are applicable, we have compressed $\varphi(\mathbf{x})$ to an equivalent atomic formula $t(\mathbf{x}') = 0$. In step 5, we construct a matrix M such that $\det M = t(\mathbf{x}')$ using Proposition 6. ■

Corollary 12 *Let F be one of the fields $\mathbb{Q}$ or $\mathbb{R}$. The decision problem* `SING` *for $S = F$ and $E = \{0, 1\}$ is NP-hard.*

Proof. Immediate from Lemma 11 and Proposition 8. ■

11 Lower bound for minrank over a field

We have just proven for the specific fields $GF(q)$, $\mathbb{Q}$ and $\mathbb{R}$ that the decision problem `SING` is as hard as deciding the corresponding existential first order theory. It is unlikely that this result can be generalized to an arbitrary field, since we have found a random polynomial-time algorithm for `SING` over $\mathbb{C}$ and the existential first-order theory is *NP*-hard over any field, in particular over $\mathbb{C}$. However, only one step in the proofs of Lemmas 9 and 11 does not seem to generalize to an arbitrary field — namely the reduction of a system (conjunction) of equations to a single equation, which is necessary for encoding a general existential sentence as a singularity problem. However, we observe that a *system* of equations can be encoded as a single minrank problem. In this section, we show that over any field the more general decision problem `MINRANK` is indeed as hard as the corresponding existential first order theory. Our construction will also lead to an alternative proof for the hardness of the singularity problem over the fields $GF(q)$, $\mathbb{Q}$ and $\mathbb{R}$.

Lemma 13 *Let F be a field.*

Given an existential sentence $\exists x_1 \cdots \exists x_t \,.\, \varphi(x_1, \ldots, x_t)$*, of length* m*, we can in time* $m^{O(1)}$ *construct an equivalent existential sentence* $\exists x_1 \cdots \exists x_{t'} \,.\, \psi(x_1, \ldots, x_{t'})$ *such that* ψ *contains neither negation nor disjunction; i.e.,* ψ *is a conjunction of atomic formulas,*

$$\psi(\mathbf{x}') \quad \equiv \quad p_1(\mathbf{x}') = 0 \ \wedge \cdots \wedge \ p_r(\mathbf{x}') = 0$$

for some arithmetic formulas p_i*,* $i = 1, \ldots r$*, and*

$$F \models \exists \mathbf{x}.\ \varphi(\mathbf{x}) \quad \textit{iff} \quad F \models \exists \mathbf{x}'.\ \psi(\mathbf{x}').$$

Proof. First we remove all negations from φ, using the rewriting rules of step 1 and 2 in Table 7, which are valid in any field.

Without loss of generality, we may therefore assume that we are given the existential sentence

$$\exists x_1 \cdots \exists x_t \,.\, \varphi(x_1, \ldots, x_t)$$

where φ is an unquantified formula *without* negations using variables $x_1, \ldots, x_t$.

Let φ have s subformulas $f_1, \ldots, f_s$, each of which may be atomic or composite. For each such subformula f_i, we introduce a new (existentially quantified) variable z_i, and we construct a new formula f_i' that is either atomic or the conjunction of two atomic formulas. The f_i's will be constructed such that

$$\begin{array}{l} \exists x_1 \cdots \exists x_t.\ \text{“}f_i \text{ is satisfied”} \\ \qquad\qquad \Updownarrow \\ \exists x_1 \cdots \exists x_t \exists z_1 \cdots \exists z_t.\ \text{“}z_i = 0 \text{ and } f_j' \text{ is satisfied} \\ \qquad \text{for all subformulas } f_j \text{ of } f_i \text{ (including } f_i)\text{”}. \end{array} \tag{7}$$

If f_1 is the subformula corresponding to the entire formula φ, the implications yield

$$\begin{array}{l} \exists \mathbf{x} \,.\, \varphi(\mathbf{x}) \\ \qquad\qquad \Updownarrow \\ \exists \mathbf{x}, \mathbf{z}.\ \ z_1 = 0 \ \wedge \ f_1'(\mathbf{x}, \mathbf{z}) \ \wedge \cdots \wedge \ f_s'(\mathbf{x}, \mathbf{z}). \end{array}$$

For each original subformula f_i the new formula f_i' is constructed as described in Table 8. By induction in the structure of f_i, one may verify that this construction does satisfy (7), from which the theorem follows. ■

f_i	f_i'
$p_i(\mathbf{x}) = 0$	$p_i(\mathbf{x}) = z_i$
$f_j \vee f_k$	$z_j \cdot z_k = z_i$
$f_j \wedge f_k$	$z_j \cdot z_k = z_i \ \wedge \ z_j + z_k = z_i$

Table 8: Subconstruction for elimination of $\vee$.

Lemma 14 *Let F be a field.*

Given an existential sentence φ of length m, we can in time $n^{O(1)}$ construct an integer k and an $n \times n$ matrix with entries from $\{0,1\} \cup \{x_1, x_2, \ldots, x_t\}$, where $n = O(m)$ such that

$$\text{minrank}_F(M) \leq k \quad \textit{iff} \quad F \models \varphi.$$

Proof. Let an existential sentence be given. First we remove all negations and disjunctions using the construction of Lemma 13.

Without loss of generality, we may therefore assume that we are given the existential sentence

$$\exists \mathbf{x}.\ p_1(\mathbf{x}) = 0 \ \wedge \cdots \wedge \ p_r(\mathbf{x}) = 0$$

for some arithmetic formulas p_i, $i = 1, \ldots r$.

By Proposition 6, we may for each $p_i(x_1, \ldots, x_t)$ find an $n_i \times n_i$ matrix M_i with entries from $\{0,1\} \cup \{x_1, x_2, \ldots, x_t\}$ such that $\det M_i = p_i(x_1, \ldots, x_t)$ and $\text{minrank}_F(M_i) \geq n_i - 1$.

Let $n = \sum_{i=1}^{r} n_i$, let $k = \sum_{i=1}^{r} (n_i - 1)$, and construct the $n \times n$ matrix M by placing $M_1, \ldots, M_r$ consecutively on the main diagonal and zeroes elsewhere. Clearly, $\text{minrank}_F(M) \geq k$ and rank $M = k$ only when all the polynomials p_i are simultaneously zero; therefore $\text{minrank}_F(M) \leq k$ iff $F \models \varphi$. ■

Corollary 15 *Let F be a field. The decision problem* `MINRANK` *for $S = F$ and $E = \{0,1\}$ is NP-hard.*

Proof. Immediate from Lemma 14 and Proposition 8. ■

Lemma 13 can also be used to give alternative proofs for Lemmas 9 and 11.

Lemma 16 *Let an existential sentence* $\exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t)$ *of length* m *be given.*

If $F = GF(q)$ *is a fixed finite field, then we can in time* $n^{O(1)}$ *construct two* $n \times n$ *matrices* M' *and* M'' *with entries from* $\{0,1\} \cup \{x_1, \ldots, x_t\}$*, where* $n = O(mq)$ *such that*

$$F \models \exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t) \quad \textit{iff} \quad \exists (a_1, \ldots, a_t) \in F^t.\ \det M'(a_1, \ldots, a_t) = 0$$

and

$$F \models \exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t) \quad \textit{iff} \quad \exists (a_1, \ldots, a_t) \in F^t.\ \det M''(a_1, \ldots, a_t) \neq 0.$$

If F *is one of the fields* $\mathbb{Q}$ *and* $\mathbb{R}$ *then we can in time* $n^{O(1)}$ *construct an* $n \times n$ *matrix* M *with entries from* $\{0,1\} \cup \{x_1, \ldots, x_t\}$*, where* $n = O(m)$ *such that*

$$F \models \exists x_1 \cdots \exists x_t.\ \varphi(x_1, \ldots x_t) \quad \textit{iff} \quad \exists (a_1, \ldots, a_t) \in F^t.\ \det M(a_1, \ldots, a_t) = 0.$$

Proof. Let an existential sentence be given. First we remove all negations and disjunctions using the construction of Lemma 13. Without loss of generality, we may therefore assume that we are given the existential sentence

$$\exists \mathbf{x}.\ p_1(\mathbf{x}) = 0 \ \wedge \cdots \wedge \ p_r(\mathbf{x}) = 0$$

for some arithmetic formulas p_i, $i = 1, \ldots r$.

If F is the finite field $GF(q)$, we use the property that the function $x \mapsto x^{q-1}$ maps 0 to 0 and maps any nonzero number to 1; i.e.,

$$\begin{array}{lcl} GF(q) & \models & \exists \mathbf{x}.\ p_1(\mathbf{x}) = 0 \ \wedge \cdots \wedge \ p_r(\mathbf{x}) = 0 \\ \Updownarrow & & \\ GF(q) & \models & \exists \mathbf{x}.\ 1 - (1 - p_1(\mathbf{x})^{q-1}) \cdot \ldots \cdot (1 - p_r(\mathbf{x})^{q-1}) = 0 \\ \Updownarrow & & \\ GF(q) & \models & \exists \mathbf{x}.\ (1 - p_1(\mathbf{x})^{q-1}) \cdot \ldots \cdot (1 - p_r(\mathbf{x})^{q-1}) \neq 0. \end{array}$$

If F is one of the fields $\mathbb{Q}$ and $\mathbb{R}$, we use the property that the function $x \mapsto x^2$ maps 0 to 0 and maps any nonzero number to a positive number, i.e.,

$$\begin{array}{lcl} F & \models & \exists \mathbf{x}.\ p_1(\mathbf{x}) = 0 \ \wedge \cdots \wedge \ p_r(\mathbf{x}) = 0 \\ \Updownarrow & & \\ F & \models & \exists \mathbf{x}.\ p_1(\mathbf{x})^2 + \cdots + p_r(\mathbf{x})^2 = 0. \end{array}$$

One may now construct the matrices with the postulated properties using Lemma 6. ■

12 Upper bounds for minrank over a field

In this section, we prove that the minrank problem over a field is no harder than deciding the corresponding existential first order theory. Combined with our earlier results, this implies that the decision problem `MINRANK` is in fact equivalent (under polynomial-time transformations) to deciding the corresponding existential first-order theory. In addition we inherit the upper bounds of Table 4.

We start by giving the reduction for matrices that use only constants 0 and 1, and afterwards extend the result to more general constants.

Lemma 17 *Let F be a field.*

Given an $n \times n$ matrix M with entries from $\{0,1\} \cup \{x_1, x_2, \ldots, x_t\}$, and some $k \leq n$, we may in time $n^{O(1)}$ construct an existential sentence φ such that

$$\text{minrank}_F(M) \leq k \quad \textit{iff} \quad F \models \varphi.$$

Proof. Given $(n \times n)$ matrix M with variables $x_1, x_2, ..., x_t$ and constants from $\{0,1\}$, we express (in a first-order existential sentence) the assertion that some k columns of M span all columns of M. For this purpose we introduce n new variables $y_1, y_2, \ldots, y_n$ in addition to the variables already occurring in M. Define the modified matrixM', where $M'_{ij} = y_j \cdot M_{ij}$; i.e., each column of M' is a multiple (possibly zero) of the corresponding column in M. We also introduce n^2 new variables $z_{11}, \ldots, z_{nn}$ forming an $n \times n$ matrix Z. The assertion minrank$(M) \leq k$ is now equivalent to the following assertion: it is possible to choose the y_j's and z_{ij}'s in such a way that at most k of the y_j's are nonzero and the matrix equation $M' \cdot Z = M$ holds.

Our sentence will be an existential quantification of a conjunction of two formulas. The first one f_1 will assert that at most k of the y_j's are nonzero, and the second one f_2 will assert that the matrix equation $M' \cdot Z = M$ holds.

Construction of f_1: We use the elementary symmetric functions defined by

$$\sigma_j(y_1, \ldots, y_n) = \sum_{A \subseteq \{1,\ldots,n\} \wedge |A|=j} \prod_{i \in A} y_i,$$

for $j = 1, \ldots, n$. These functions satisfy the following property: there are at most k nonzero y_j's if and only if

$$\sigma_{k+1}(y_1, \ldots, y_n) = 0 \ \wedge \ \sigma_{k+2}(y_1, \ldots, y_n) = 0 \ \wedge \ \cdots \ \wedge \ \sigma_n(y_1, \ldots, y_n) = 0.$$

The "only if" direction is trivially satisfied. For the "if" direction one can prove the result by induction: For the basis of the induction consider $\sigma_n(y_1, \ldots, y_n) = \prod_{i=1}^{n} y_i$. If $\sigma_n(y_1, \ldots, y_n) = 0$ then, since F contains no zero divisors, we must have some $y_i = 0$, and without loss of generality assume $y_n = 0$. Since $y_n = 0$, $\sigma_{n-1}(y_1, \ldots, y_n)$ reduces to $\prod_{i=1}^{n-1} y_i$, and the argument can be repeated to prove that in total at most k of the y_j's are nonzero.

We need to find a short formula expressing that $\sigma_j(y_1, \ldots, y_n) = 0$. Consider the polynomial $p(z, y_1, y_2, \ldots, y_n) = (z + y_1)(z + y_2) \cdots (z + y_n) = z^n + \sigma_1(y_1, \ldots, y_n) z^{n-1} + \cdots + \sigma_n(y_1, \ldots, y_n)$. Using this equality, the school method for multiplying out polynomials gives an arithmetic circuit of size $O(n^3)$ that computes $\sigma_j(y_1, \ldots, y_n)$ for all $j = 1, \ldots, n$ simultaneously. This circuit can be understood as a straight-line program of length $s = O(n^3)$ using the operations $\{+, \cdot\}$. Let the atomic formula h_i be $w = u + v$ (respectively $w = u \cdot v$) if the i'th line of the straight-line program is $w \leftarrow u + v$ (respectively $w \leftarrow u \cdot v$). Without loss of generality, we may assume that the variables used in the straight-line program are $w_1, \ldots, w_s$ in addition to input variables $y_1, \ldots, y_n$ and output variables $s_1, \ldots, s_n$ computing $\sigma_1, \ldots, \sigma_n$. A suitable formula for f_1 is

$$f_1 \quad \equiv \quad h_1 \ \wedge \ \cdots \ \wedge \ h_s \ \wedge \ s_{k+1} = 0 \ \wedge \ \cdots \ \wedge \ s_n = 0.$$

Construction of f_2: We need to express that $M' \cdot Z = M$. The ij'th entry in the matrix product is $\sum_{k=1}^{n} y_k M_{ik} z_{kj}$. Therefore construct the atomic formula

$$g_{ij} \quad \equiv \quad y_1 M_{i1} z_{1j} + y_2 M_{i2} z_{2j} + \cdots + y_n M_{in} z_{nj} = M_{ij},$$

and define

$$f_2 \quad \equiv \quad g_{11} \ \wedge \ g_{12} \ \wedge \ \cdots \ \wedge \ g_{nn}.$$

The sentence required by the lemma is thus

$$\begin{aligned} &\mathrm{minrank}_F(M) \leq k \quad \equiv \\ &\quad \exists x_1 \cdots \exists x_t \ \exists y_1 \cdots \exists y_n \ \exists z_{11} \cdots \exists z_{nn} \ \exists w_1 \cdots \exists w_s \ \exists s_1 \cdots \exists s_n \ . \ f_1 \ \wedge \ f_2. \end{aligned}$$

■

We restricted the constants in our existential sentences to 0 and 1 in order to apply the upper bounds of Table 4. However, an analogue of Lemma 17

does actually hold for the minrank problem over matrices containing algebraic constants, because algebraic constants can be defined by short first-order sentences.

- Over any field, the constant 2 is defined by

$$\varphi(x) \quad \equiv \quad x = 1 + 1.$$

- Over a field with characteristic different from 2, the constant $-\frac{3}{2}$ is defined by

$$\varphi(x) \quad \equiv \quad x \cdot (1 + 1) + 1 + 1 + 1 = 0.$$

- Over $\mathbb{R}$, the constant $\sqrt{2}$ is defined by

$$\varphi(x) \quad \equiv \quad \exists y.\ x \cdot x = 1 + 1 \ \wedge \ y \cdot y = x$$

  (The last part ensures that we get the positive of the two square roots.)

- Over any field, the constant 15 is defined by

$$\varphi(x) \quad \equiv \quad \exists y \exists z \exists w.\ x = 1+y+z+w \ \wedge \ y = 1+1 \ \wedge \ z = y+y \ \wedge \ w = z+z$$

  (We use a repeated doubling strategy to make the defining formula have length proportional to the usual binary representation of the integer 15.)

- Over $\mathbb{C}$, the constants i and $-i$ are defined by

$$\varphi(x, y) \quad \equiv \quad x \cdot x + 1 = 0 \ \wedge \ y \cdot y + 1 = 0 \ \wedge \ x + y = 0$$

  (Note that i and $-i$ can not be defined separately, since i alone can only be defined up to conjugation, the only nontrivial isomorphism on $\mathbb{C}$.)

If F is a field, define its *prime field* to be the intersection of all subfields of F [8, §V.5]. Clearly, the prime field underlying $\mathbb{C}$ and $\mathbb{R}$ is $\mathbb{Q}$, and $GF(q)$ is a finite-dimensional algebraic extension of its underlying prime field (which is $GF(p)$ for some prime p). For a field F let A_F be the set of all numbers that are algebraic over the prime field underlying F.

Proposition 18 *Let P be a prime field. Let $\{e_1, \ldots, e_t\} \subseteq A_P$. Let F be the smallest extension field containing all the constants $\{e_1, \ldots, e_t\}$. Let a standard representation of F as a k-dimensional vector space over P (with vector arithmetic defined using an irreducible polynomial) be given. Let the representation of the constants $\{e_1, \ldots, e_t\}$ as vectors of binary numbers be given.*

It is possible to construct an existential first order formula $\varphi(x_1, \ldots, x_t)$ defining $\{e_1, \ldots, e_t\}$ in time polynomial in the combined bit length of all the constant representations.

Proof. Left to the reader. ■

The generalization of Lemma 17 is the following.

Lemma 19 *Let F be a field. Let F' be a finite dimensional algebraic extension of the prime field underlying F. Let $E \subseteq F' (\subseteq A_F)$.*

Given an $n \times n$ matrix M with entries from $E \cup \{x_1, x_2, \ldots, x_t\}$, and some $k \leq n$, we may in time $(ns)^{O(1)}$ construct an existential sentence φ such that

$$\text{minrank}_F(M) \leq k \quad \textit{iff} \quad F \models \varphi,$$

where s denotes the maximum bit length of the representation of an entry in M (using binary numbers/quotients for prime field elements and vectors of these for algebraic numbers).

Proof. Use the construction from the proof of Lemma 17 combined with the construction of Proposition 18. ■

Corollary 20 *Let F be a field. Let F' be a finite dimensional algebraic extension of the prime field underlying F. Let $S = F$ and let $\{0, 1\} \subseteq E \subseteq F'$.*

The decision problem `MINRANK` *is equivalent (under polynomial-time transformations) to deciding $ETh(F)$.*

If F is one of the fields $\mathbb{Q}$ or $\mathbb{R}$, then the decision problems `SING` *and* `MINRANK` *are equivalent by polynomial-time transformation.*

If F is a fixed p-adic field $\mathbb{Q}_p$, then the decision problem `MINRANK` *is solvable in* doubly exponential space.

If F is one of the fields $\mathbb{R}$ and $\mathbb{C}$ then the decision problem `MINRANK` *is in PSPACE.*

Proof. Immediate from Lemmas 19, 14, 11 and the bounds cited in Table 4. ■

13 Tight approximation of minrank is *NP*-hard

In this section, we consider the following approximation problem (parametrized with $\varepsilon > 0$) associated with the minrank problem.

$(1+\varepsilon)$-APXMINRANK

Let R be a commutative ring. Let $E, S \subseteq R$.

Input: a matrix $M = M(x_1, \ldots, x_t)$ with entries in $E \cup \{x_1, \ldots, x_r\}$.

Output: some $a_1, \ldots, a_t \in S$ such that

$$\text{rank } M(a_1, \ldots, a_t) \leq (1+\varepsilon) \cdot \text{minrank}_S(M).$$

We prove that $(1+\varepsilon)$-APXMINRANK is *NP*-hard for ε sufficiently small, when R is $\mathbb{Z}$ or a field. The tool will be reduction from the approximation version of EXACT3SAT. Consider the following problem.

$(1-\varepsilon)$-MAXEXACT3SAT

Input: a conjunction of clauses $C = C_1 \wedge \cdots \wedge C_k$, where each clause contains exactly three distinct literals $C_i = (l_{i1} \vee l_{i2} \vee l_{i3})$, and each literal is one of the Boolean variables $\{y_1, \ldots, y_r\}$ or its negation.

For $(b_1, \ldots, b_r) \in \{0,1\}^r$, let $\text{numb}(C, b_1, \ldots, b_r)$ be the number of clauses in C that are satisfied under the assignment $y_i \mapsto b_i$, and let

$$\text{maxnumb}(C) = \max_{(b_1,\ldots,b_r)\in\{0,1\}^r} \text{numb}(C, b_1, \ldots, b_r).$$

Output: some truth assignment $b_1, b_2, \ldots, b_r \in \{0,1\}$ such that

$$\text{numb}(C, b_1, \ldots, b_r) \geq (1-\varepsilon) \cdot \text{maxnumb}(C).$$

Proposition 21 *For $\varepsilon < \frac{1}{27}$ there is no polynomial-time algorithm for* $(1-\varepsilon)$-MAXEXACT3SAT *unless* $P = NP$.

Proof. See Bellare et al. [1, pages 48ff]. ■

To prove the non-approximability of minrank, we need a special type of reduction first defined by Papadimitriou and Yannakakis [13]. Since we only use the reduction in a single case, we specialize the definition to the concrete application.

Given $E, S \subseteq R$, MAXEXACT3SAT is said to *L-reduce* to APXMINRANK with parameters α, β, if there exist two polynomial time computable functions f and g such that for a given instance C of MAXEXACT3SAT,

1. Algorithm f produces matrix M with entries in $E \cup \{x_1, \ldots, x_t\}$ such that
$$\text{minrank}_S(M) \leq \alpha \cdot \text{maxnumb}(C);$$
2. Given any substitution $(a_1, a_2, \ldots, a_t) \in S^t$ for the variables in M, g produces a truth assignment $(b_1, b_2, \ldots, b_r) \in \{0,1\}^r$ such that
$$|\text{maxnumb}(C) - \text{numb}(C, b_1, b_2, \ldots, b_r)| \leq \beta \cdot |\text{minrank}_S(M) - \text{rank } M(a_1, a_2, \ldots, a_t)|.$$

L-reduction preserves approximability.

Proposition 22 *Let $E, S \subseteq R$ be given. If* MAXEXACT3SAT *L-reduces to* APXMINRANK *with parameters $\alpha, \beta \geq 0$ and* $(1+\varepsilon)$-APXMINRANK *has a polynomial time solution then* $(1-\alpha\beta\varepsilon)$-MAXEXACT3SAT *has a polynomial time solution.*

Proof. The polynomial time solution for $(1-\alpha\beta\varepsilon)$-MAXEXACT3SAT works as follows. Given an instance $C(y_1, \ldots, y_r)$ of MAXEXACT3SAT, compute an instance $M(x_1, \ldots, x_t)$ of APXMINRANK using the function f. Find a substitution $(a_1, \ldots, a_t)$ for $(x_1, \ldots, x_t)$ using the polynomial time solution for $(1+\varepsilon)$-APXMINRANK, and transform this substitution into a truth assignment $(b_1, \ldots, b_r)$ for $(y_1, \ldots, y_r)$ using the function g. We verify the $(1-\alpha\beta\varepsilon)$ bound by a computation:

$$
\begin{aligned}
&|\text{maxnumb}(C) - \text{numb}(C, b_1, \ldots, b_r)| \\
&\quad \leq \beta \cdot |\text{minrank}_S(M) - \text{rank } M(a_1, \ldots, a_t)| \\
&\quad \leq \beta\varepsilon \cdot \text{minrank}_S M \\
&\quad \leq \alpha\beta\varepsilon \cdot \text{maxnumb}(C).
\end{aligned}
$$

■

Lemma 23 *Let R be a commutative ring without zero divisors, and let $\{0,1\} \subseteq S \subseteq R$ and $E = \{0,1\}$.* `MAXEXACT3SAT` *L-reduces to* `APXMINRANK` *with parameters $\alpha = \frac{65}{7}$ and $\beta = 1$.*

Proof. First, we describe the function f. Assume we have an instance of `MAXEXACT3SAT`, viz. a conjunction of clauses $C = C_1 \wedge \cdots \wedge C_k$, where each clause contains three distinct literals $C_i = (l_{i1} \vee l_{i2} \vee l_{i3})$, and each literal is one of the Boolean variables $\{y_1, \ldots, y_r\}$ or its negation.

For each clause C_i, there will be a 12×12 matrix M_i, containing four smaller 3×3 matrices down the diagonal and zeroes elsewhere. The four smaller matrices are one for each of the three variables occurring in the clause and one for the clause itself.

Each Boolean variable y_j is represented by two arithmetic variables x_{j1} and x_{j2}. The variable x_{j1} being zero represents y_j being true, and x_{j2} being zero represents y_j being false. We can ensure that not both of x_{j1} and x_{j2} are zero by requiring

$$x_{j1} + x_{j2} = 1 \tag{8}$$

We allow the case that neither x_{j1} nor x_{j2} is zero.

For each of the three variables occurring in a clause, there will be a matrix ensuring (8); i.e., for $s = 1, 2, 3$, if $l_{is} = y_j$ or $l_{is} = \overline{y_j}$ then

$$A_{is} = \begin{bmatrix} 1 & x_{j1} & x_{j2} \\ 1 & 1 & 0 \\ 1 & 0 & 1 \end{bmatrix}$$

The matrix A_{is} always has rank at least 2, and has rank exactly 2 when (8) is satisfied, since $\det A_{is} = 1 - x_{j1} - x_{j2}$.

If $C_i = (y_{j_1} \vee y_{j_2} \vee y_{j_3})$, the fourth matrix will be

$$B_i = \begin{bmatrix} x_{j_1 1} & 1 & 0 \\ 0 & x_{j_2 1} & 1 \\ 0 & 0 & x_{j_3 1} \end{bmatrix}$$

(If $\overline{y_j}$ occurs in C_i instead of y_j, then replace x_{j1} with x_{j2} in matrix B_i.)

The matrix B_i always has rank at least 2, and has rank exactly 2 when $x_{j_1 1} = 0$ or $x_{j_2 1} = 0$ or $x_{j_3 1} = 0$.

The function f returns the matrix

$$M = \mathrm{diag}(M_1, \ldots, M_k), \quad \text{where} \quad M_i = \mathrm{diag}(A_{i1}, A_{i2}, A_{i3}, B_i).$$

Clearly, f can be computed in polynomial time.

Clearly, $\mathrm{minrank}_S(M) \leq k \cdot (4 \cdot 2) + (k - \mathrm{maxnumb}(C)) = 9k - \mathrm{maxnumb}(C)$. We know that $\mathrm{maxnumb}(C) \geq \frac{7k}{8}$, since the expected fraction of true clauses using a random truth assignment is at least $\frac{7}{8}$. Combining, we get that

$$\begin{aligned} \mathrm{minrank}_S(M) &\leq 9k - \mathrm{maxnumb}(C) \\ &\leq 9 \cdot \frac{8}{7} \mathrm{maxnumb}(C) - \mathrm{maxnumb}(C) \\ &= \frac{65}{7} \mathrm{maxnumb}(C), \end{aligned}$$

which proves the assertion about α.

We still need to describe the function g. Let a substitution $a_{11}, a_{12}, \ldots, a_{r1}, a_{r2} \in S^{2r}$ for the arithmetic variables in M be given. Construct a truth assignment $b_1, \ldots, b_r$ for the Boolean variables in C as follows. If $a_{j1} = 0$, then let $b_j = 1$, otherwise if $a_{j2} = 0$ then let $b_j = 0$, but if both $a_{j1} \neq 0$ and $a_{j2} \neq 0$ then let b_j take an arbitrary value. Clearly, g can be computed in polynomial time.

If clause C_i is not satisfied under the truth assignment $b_1, \ldots, b_r$, then matrix M_i will have rank at least 9 under the substitution $a_{11}, a_{12}, \ldots, a_{r1}, a_{r2}$, because either $a_{j1} = a_{j2} = 0$ for some variable y_j occurring in C_i and then one of A_{is} will have rank 3, or matrix B_i will have rank 3.

Therefore, $k - \mathrm{numb}(C, b_1, \ldots, b_r) \leq \mathrm{rank}\ M(a_{11}, a_{12}, \ldots, a_{r1}, a_{r2}) - 8k$, which combined with our earlier inequality, $\mathrm{minrank}_S(M) \leq 9k - \mathrm{maxnumb}(C)$, implies

$$\begin{aligned} &\mathrm{maxnumb}(C) - \mathrm{numb}(C, b_1, \ldots, b_r) \\ &\quad \leq 9k - \mathrm{minrank}_S(M) + \mathrm{rank}\ M(a_{11}, a_{12}, \ldots, a_{r1}, a_{r2}) - k - 8k \\ &\quad = \mathrm{rank}\ M(a_{11}, a_{12}, \ldots, a_{r1}, a_{r2}) - \mathrm{minrank}_S(M), \end{aligned}$$

which proves the assertion about β. ■

Theorem 24 *Let R be a commutative ring without zero divisors, and let $\{0,1\} \subseteq S \subseteq R$ and $E = \{0,1\}$. For $\varepsilon < \frac{7}{1755} \approx .0039886$ there is no polynomial time solution for $(1+\varepsilon)$-*APXMINRANK *unless $P = NP$.*

Proof. Combine Propositions 21 and 22 with Lemma 23. ■

14 The case when each variable occurs exactly once

In previous sections we have been considering matrices $M = M(x_1, x_2, \ldots, x_t)$ with entries in $E \cup \{x_1, x_2, \ldots, x_t\}$, and each variable can occur arbitrarily often in M. In this section and the next, we restrict our attention to matrices where each variable occurs exactly once, and we call such matrices *eveo*.

Definition. A polynomial $p(x_1, x_2, \ldots, x_t)$ is said to be *multi-affine* over a field F if, for every substitution of all variables but one (say x_i) with field elements, the result can be expressed in the form $ax_i + b$ with $a, b \in F$.

Alternatively, p is multi-affine if every variable occurs with degree 0 or 1 in every term. For example, $2xyz + 3z + 4x + 5$ is multi-affine over $\mathbb{Q}$. Note that the determinant of an eveo matrix is multi-affine.

The following lemmas will prove useful. We say that a polynomial $p(x_1, x_2, \ldots, x_t)$ is *identically zero* over a field F if $p(a_1, a_2, \ldots, a_t) = 0$ for all $a_1, a_2, \ldots, a_t \in F$.

Lemma 25 *Let p be a multi-affine polynomial over a field F. Then p is identically zero over F iff p is the zero polynomial.*

Note that this theorem is *not* necessarily true for polynomials in which variables occur with higher degree; for example, the polynomial $x^2 - x$ is not the zero polynomial, but is identically zero over $GF(2)$.

Proof. If p is the zero polynomial, the result is evident.

Now assume p is not the zero polynomial. We will prove by induction on the number of variables that p is not identically zero. If $t = 1$, then $p(x) = ax + b$, and at least one of a, b is nonzero. If b is nonzero, then we

can set $x = 0$ to get a nonzero value. If a is nonzero, then set $x = (1 - b)/a$ to get the value 1.

Now assume the result is true for $t < k$ variables; we prove it for $t = k$ variables. Let $p = qx + r$, where q, r are multi-affine polynomials in $k - 1$ variables. Then since p is not the zero polynomial, either q or r must be different from the zero polynomial. By induction. either q or r takes a nonzero value. Substitute values for the $k - 1$ variables to obtain $ax + b$, where not both a and b are zero. Then, as above, $ax + b$ takes a non-zero value in F. ■

Corollary 26 *A multi-affine polynomial is identically zero over a field F iff it is identically zero over some extension field $F' \supseteq F$.*

Lemma 27 *A multi-affine function over a field is either constant or takes all values in the field.*

Proof. The proof is by induction on d, the number of variables. If $d = 1$, then $p(x) = ax + b$. If p is non-constant, then $a \neq 0$. Then to get $p(x) = c$, choose $x = (c - b)/a$.

Otherwise, p is a function of $t \geq 2$ variables. Choose any variable that occurs at least once, say x. Write $p = ax + b$, where a, b are multi-affine polynomials in $t - 1$ variables. The polynomials a and b cannot be both constant. If a is a constant, choose any assignment of variables to b, forcing b to take the value b'; now set $x = (b' - c)/a$. If a is non-constant, then by induction it takes on all values in F, so choose an assignment to the variables in a that makes it nonzero; this can be done by Lemma 25. This assignment of variables gives a the value a' and b the value b', and now set $x = (b' - c)/a'$. ■

Theorem 28 *For all fields F, and all eveo matrices M, we can compute* $\text{maxrank}_F(M)$ *in random polynomial time.*

Proof. We mimic the proof of Theorem 2. Let M be an $n \times n$ eveo matrix. If the field F has at least $2n$ elements, then the proof goes through essentially unchanged, with V any subset of F of cardinality $2n$. Otherwise, choose an appropriate field extension F' with at least $2n$ elements. By Corollary 26 a

minor is not identically zero over F' iff it is not identically zero over F, so we may compute maxrank over F' instead of over F. ■

Now recall the singularity problem.

Theorem 29 *If F is a field, and M is an eveo matrix, then the decision problem* SING *is in the complexity class RP.*

Proof. By Lemmas 25 and 27, it is enough to ensure that the determinant $\det M$ is not a nonzero constant polynomial. Mimic the proof of Theorem 5, using Corollary 26, if necessary, to extend the base field. ■

15 The minrank problem for row-partitionable matrices

In this section we show that the minrank problem is solvable in deterministic polynomial time if the matrix has a certain special form, in which each variable appears only once and there is a division between the variable and non-variable entries.

More formally, let M be an $m \times n$ matrix with entries chosen from $E \cup \{x_1, x_2, \ldots, x_t\}$. We say that M is *row-partitionable* if

(a) each variable x_i occurs exactly once in M; and

(b) for each row i there exists an index k_i such that $a_{ij} \in E$ if $1 \leq j \leq k_i$, and $a_{ij} \notin E$ if $k_i < j \leq n$.

As an example, the following matrix is row-partitionable:

$$M = \begin{bmatrix} 3 & 7 & -2 & x_1 & x_2 \\ 2 & 4 & x_3 & x_4 & x_5 \\ -3 & 5 & 6 & 2 & x_6 \\ 7 & 2 & 9 & 1 & 4 \end{bmatrix}$$

The main motivation for this subproblem comes from the theory of *rational series*; for an introduction to this area, see [2]. Let f be a formal power series in noncommuting variables over a field F. Then f is said to be *rational* if it can be expressed using the operations sum, product, and quasi-inverse

(the map sending $x \to 1/(1-x)$). The series f is said to be *recognizable* if the coefficient of the term corresponding to w (which is written as (f, w)) can be computed as follows: there is a matrix-valued homomorphism μ, a row matrix λ, and a column matrix γ such that $(f, w) = \lambda\mu(w)\gamma$. A well-known theorem due to Schützenberger (e.g., [2, Thm. 6.1]) proves that a formal power series is rational iff it is recognizable. In this case the dimension of the smallest possible matrix representation (the dimension of the square matrix $\gamma\lambda$) is an invariant called the *rank* of the rational series. The following problem now arises [7, 16]: given a (not necessarily rational) formal power series f, compute the smallest possible rank $R_f(n)$ of any rational series agreeing with f on all terms of total degree at most n.

It can be shown that this number $R_f(n)$ is equal to the minrank of an associated Hankel-like matrix $M(f, n)$. More specifically, we have $R_f(n) = \text{minrank}_F(M(f, n))$, where the rows of $M(f, n)$ are labeled with words w of length $\leq n$, the columns are labeled with words x of length $\leq n$, and the entry in the row corresponding to w and the column corresponding to x is (f, wx) if $|wx| \leq n$, and a unique indeterminate otherwise. It is easy to see that this particular $M(f, n)$ is row-partitionable.

Consider the following algorithm.

MR$(M = (a_{ij})_{1 \leq i \leq m, 1 \leq j \leq n})$
(1) rearrange rows so that $k_1 \geq k_2 \geq \cdots \geq k_m$;
(2) if there exists u, $1 \leq u \leq k_1$ such that $a_{1u} \neq 0$,
set $r \leftarrow 1$; $T \leftarrow \{1\}$
else
set $r \leftarrow 0$; $T \leftarrow \emptyset$
(3) for $s = 2$ to m do
if the vector $(a_{s1}, a_{s2}, \ldots, a_{s,k_s})$ is not linearly dependent on $(a_{ij})_{i \in T,\ 1 \leq j \leq k_s}$
set $r \leftarrow r + 1$; $T \leftarrow T \cup \{s\}$
(4) return(r)

Theorem 30 *Let F be a field. Then algorithm* **MR** *correctly computes* $\text{minrank}_F(M)$ *and uses $O(m^3 n)$ field operations.*

To prove correctness, we first observe that the reordering in step (1) cannot change $\text{minrank}_F(M)$.

Next, we observe that the following invariants hold before the loop step corresponding to s is performed:

(a) for all possible assignments to the variables, the rows in the set T are linearly independent;

(b) for each assignment to the variables in the rows of T, there exists an assignment to the variables in the rows $\overline{T} = \{1, 2, \ldots, s-1\} - T$ such that each of the rows in $\overline{T}$ is dependent on a row of T.

These invariants clearly hold after step (2). We now prove by induction on s that they hold throughout the algorithm.

Suppose the invariants hold up to step $s-1$. At step s, we consider row s of M. If $(a_{s1}, \ldots, a_{s,k_s})$ is not dependent on $(a_{ij})_{i \in T,\ 1 \le j \le k_s}$, then for any assignment of the variables row s of M is not dependent on the rows in T, so by adding s to T we preserve part (a) of the invariant, and part (b) is unaffected. If, on the other hand, $a = (a_{s1}, \ldots, a_{s,k_s})$ is dependent on $M' = (a_{ij})_{i \in T,\ 1 \le j \le k_s}$, then write a as a linear combination of the rows of M'. We can then assign the variables in row s of M appropriately so that the entire row s is a linear combination of the rows of T. Then part (b) of the invariant is preserved, and part (a) is unaffected. This completes the proof of correctness.

To complete the proof of the theorem, it suffices to observe that we can test to see if row s is dependent on rows of T in at most $O(m^2 n)$ field operations, and this step is performed at most m times. ■

References

[1] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs and non-approximability — towards tight results (3rd revision). Report Series 1995, Revision 02 of ECCC TR95-024, Electronic Colloqium on Computational Complexity, `http://www.eccc.uni-trier.de/eccc/`, December 1995. Earlier results appeared in *Proc. 36th Ann. Symp. Found. Comput. Sci.* (1995), 422–431.

[2] J. Berstel and C. Reutenauer. *Rational Series and Their Languages*, Vol. 12 of *EATCS Monographs on Theoretical Computer Science*. Springer-Verlag, 1988.

[3] J. Canny. Some algebraic and geometric computations in PSPACE. In *Proc. Twentieth ACM Symp. Theor. Comput.*, pp. 460–467, 1988.

[4] L. Egidi. The complexity of the theory of p-adic numbers. In *Proc. 34th Ann. Symp. Found. Comput. Sci.*, pp. 412–421, 1993.

[5] H. B. Enderton. *A Mathematical Introduction to Logic*. Academic Press, 1972.

[6] J. Friedman. A note on matrix rigidity. *Combinatorica* **13** (1993), 235–239.

[7] C. Hespel. Approximation de séries formelles par des séries rationnelles. *RAIRO Inform. Théor.* **18** (1984), 241–258.

[8] T. W. Hungerford. *Algebra*, Vol. 73 of *Graduate Texts in Mathematics*. Springer-Verlag, 1987.

[9] D. Ierardi. Quantifier elimination in the theory of an algebraically-closed field. In *Proc. Twenty-first Ann. ACM Symp. Theor. Comput.*, pp. 138–147, 1989.

[10] S. Lang. *Algebra*. Addison-Wesley, 1971.

[11] S. V. Lokam. Spectral methods for matrix rigidity with applications to size-depth tradeoffs and communication complexity. In *Proc. 36th Ann. Symp. Found. Comput. Sci.*, pp. 6–16, 1995.

[12] Y. V. Matiyasevich. *Hilbert's Tenth Problem*. The MIT Press, 1993.

[13] C. H. Papadimitriou and M. Yannakakis. Optimization, approximation, and complexity classes. *J. Comput. System Sci.* **43** (1991), 425–440.

[14] J. Renegar. On the computational complexity and geometry of the first-order theory of the reals. part I: Introduction. preliminaries. the geometry of semi-algebraic sets. the decision problem for the existential theory of the reals. *J. Symbolic Comput.* **13** (1992), 255–299.

[15] J. T. Schwartz. Fast probabilistic algorithms for verification of polynomial identities. *J. Assoc. Comput. Mach.* **27** (1980), 701–717.

[16] J. O. Shallit. On approximation by rational series in noncommuting variables. Unpublished manuscript, in preparation, 1996.

[17] L. Valiant. Graph-theoretic arguments in low-level complexity. In *6th Mathematical Foundations of Computer Science*, Vol. 197 of *Lecture Notes in Computer Science*, pp. 162–176. Springer-Verlag, 1977.

[18] L. G. Valiant. Completeness classes in algebra. In *Proc. Eleventh Ann. ACM Symp. Theor. Comput.*, pp. 249–261, 1979.

Recent Publications in the BRICS Report Series

RS-96-33 Jonathan F. Buss, Gudmund Skovbjerg Frandsen, and Jeffrey Outlaw Shallit. *The Computational Complexity of Some Problems of Linear Algebra*. September 1996. 39 pp.

RS-96-32 P. S. Thiagarajan. *Regular Trace Event Structures*. September 1996. 34 pp.

RS-96-31 Ian Stark. *Names, Equations, Relations: Practical Ways to Reason about 'new'*. September 1996. ii+22 pp.

RS-96-30 Arne Andersson, Peter Bro Miltersen, and Mikkel Thorup. *Fusion Trees can be Implemented with AC^0 Instructions only*. September 1996. 8 pp.

RS-96-29 Lars Arge. *The I/O-Complexity of Ordered Binary-Decision Diagram Manipulation*. August 1996. 35 pp. An extended abstract version appears in Staples, Eades, Kato, and Moffat, editors, *Algorithms and Computation: 6th International Symposium*, ISAAC '95 Proceedings, LNCS 1004, 1995, pages 82–91.

RS-96-28 Lars Arge. *The Buffer Tree: A New Technique for Optimal I/O Algorithms*. August 1996. 34 pp. This report is a revised and extended version of the BRICS Report RS-94-16. An extended abstract appears in Akl, Dehne, Sack, and Santoro, editors, *Algorithms and Data Structures: 4th Workshop*, WADS '95 Proceedings, LNCS 955, 1995, pages 334–345.

RS-96-27 Devdatt Dubhashi, Volker Priebe, and Desh Ranjan. *Negative Dependence Through the FKG Inequality*. July 1996. 15 pp.

RS-96-26 Nils Klarlund and Theis Rauhe. *BDD Algortihms and Cache Misses*. July 1996. 15 pp.

RS-96-25 Devdatt Dubhashi and Desh Ranjan. *Balls and Bins: A Study in Negative Dependence*. July 1996. 27 pp.