

Generalized Wong sequences and their applications to Edmonds' problems

Gábor Ivanyos¹, Marek Karpinski², Youming Qiao³, and Miklos Santha⁴

- 1 Institute for Computer Science and Control, Hungarian Academy of Sciences, Budapest, Hungary
Gabor.Ivanyos@szttaki.mta.hu
- 2 Department of Computer Science, University of Bonn, Bonn, Germany
marek@cs.uni-bonn.de
- 3 Centre for Quantum Technologies, National University of Singapore, Singapore 117543.
cqmq@nus.edu.sg
- 4 LIAFA, Univ. Paris 7, CNRS, Paris, France / Centre for Quantum Technologies, National University of Singapore, Singapore
miklos.santha@liafa.jussieu.fr

Abstract

We design two deterministic polynomial time algorithms for variants of a problem introduced by Edmonds in 1967: determine the rank of a matrix M whose entries are homogeneous linear polynomials over the integers. Given a linear subspace \mathcal{B} of the $n \times n$ matrices over some field \mathbb{F} , we consider the following problems: *symbolic matrix rank* (SMR) is the problem to determine the maximum rank among matrices in \mathcal{B} , while *symbolic determinant identity testing* (SDIT) is the question to decide whether there exists a nonsingular matrix in \mathcal{B} . The constructive versions of these problems are asking to find a matrix of maximum rank, respectively a nonsingular matrix, if there exists one.

Our first algorithm solves the *constructive* SMR when \mathcal{B} is spanned by unknown rank one matrices, answering an open question of Gurvits. Our second algorithm solves the constructive SDIT when \mathcal{B} is spanned by triangularizable matrices, but the triangularization is not given explicitly. Both algorithms work over finite fields of size at least $n + 1$ and over the rational numbers, and the first algorithm actually solves (the non-constructive) SMR independent of the field size. Our main tool to obtain these results is to generalize Wong sequences, a classical method to deal with pairs of matrices, to the case of pairs of matrix spaces.

1998 ACM Subject Classification I.1.2 Algebraic algorithms

Keywords and phrases symbolic determinantal identity testing, Edmonds' problem, maximum rank matrix completion, derandomization, Wong sequences

Digital Object Identifier 10.4230/LIPIcs.STACS.2014.397

1 Introduction

In [8] Edmonds introduced the following problem: Given a matrix M whose entries are homogeneous linear polynomials over the integers, determine the rank of M . The problem is the same as determining the maximum rank of a matrix in a linear space of matrices over the rationals. In this paper we consider the same question and certain of its variants over more general fields.



© Gábor Ivanyos, Marek Karpinski, Youming Qiao, and Miklos Santha; licensed under Creative Commons License CC-BY

31st Symposium on Theoretical Aspects of Computer Science (STACS'14).

Editors: Ernst W. Mayr and Natacha Portier; pp. 397–408

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



SYMPOSIUM
ON THEORETICAL
ASPECTS
OF COMPUTER
SCIENCE

Let us denote by $M(n, \mathbb{F})$ the linear space of $n \times n$ matrices over a field \mathbb{F} . We call a linear subspace $\mathcal{B} \leq M(n, \mathbb{F})$ a *matrix space*. We define the *symbolic matrix rank* problem (SMR) over \mathbb{F} as follows: given $\{B_1, \dots, B_m\} \subseteq M(n, \mathbb{F})$, determine the maximum rank among matrices in $\mathcal{B} = \langle B_1, \dots, B_m \rangle$, the matrix space spanned by B_i 's. The *constructive* version of SMR is to find a matrix of maximum rank in \mathcal{B} (this is called the maximum rank matrix completion problem in [12] and in [19]). We refer to the weakening of SMR, when the question is to decide whether there exists a nonsingular matrix in \mathcal{B} , as the *symbolic determinant identity testing* problem (SDIT), the name used by [20] (in [15] this variant is called Edmonds' problem). The *constructive* version in that case is to find a nonsingular matrix, if there is one in \mathcal{B} . We will occasionally refer to any of the above problems as *Edmonds' problem*.

The complexity of the SDIT depends crucially on the size of the underlying field \mathbb{F} . When $|\mathbb{F}|$ is a constant then it is NP-hard [5], on the other hand if the field size is large enough (say $\geq 2n$) then by the Schwartz-Zippel lemma [25, 30] it admits an efficient randomized algorithm [21]. Obtaining a deterministic polynomial-time algorithm for the SDIT would be of fundamental importance, since Kabanets and Impagliazzo [20] showed that such an algorithm would imply strong circuit lower bounds which seem beyond current techniques.

Previous works on Edmonds' problems mostly dealt with the case when the given *matrices* B_1, \dots, B_m satisfy certain *property*. For example, Lovász [22] considered several cases of SMR, including when the B_i 's are of rank 1, and when they are skew symmetric matrices of rank 2. These classes were then shown to have deterministic polynomial-time algorithms [12, 23, 16, 13, 11, 19], see Section 1.1 for more details.

Another direction also studied is when instead of the given matrices, the generated *matrix space* $\mathcal{B} = \langle B_1, \dots, B_m \rangle$ satisfies certain *property*. Since such a property is just a subset of all matrix spaces, we also call it a *class* of matrix spaces. Gurvits [15] has presented an efficient deterministic algorithm for the SDIT over \mathbb{Q} , when the matrix space satisfies the so called *Edmonds-Rado property*, whose definition we shall review in Section 1.1. For now we only note that this class includes \mathbf{R}_1 , the class of *rank-1* spanned matrix spaces, where a matrix space \mathcal{B} is in \mathbf{R}_1 if and only if \mathcal{B} has a basis consisting of rank-1 matrices. This fact was first shown by Lovász [22] via a theorem of Rado and Edmonds [24, 9, 28]. Gurvits stated as an open question the complexity of the SMR for \mathbf{R}_1 over finite fields [15, page 456].

The difference between properties of matrices and properties of matrix spaces is critical for Edmonds' problems. For example, given matrices B_1, \dots, B_m , it is presumably hard¹ to determine whether $\mathcal{B} = \langle B_1, \dots, B_m \rangle$ is in \mathbf{R}_1 , and to find generating rank-1 matrices for \mathcal{B} . Thus the existence of algorithms for SMR when the B_i 's are rank-1 does not immediately imply algorithms for matrix spaces in \mathbf{R}_1 .

Our results are in line with Gurvits' work, namely we present algorithms for two classes of matrix spaces. To be specific, we consider \mathbf{R}_1 , the class of rank-1 spanned matrix spaces, and the class of (upper-)triangularizable matrix spaces, where a matrix space $\mathcal{B} \leq M(n, \mathbb{F})$ is *triangularizable* if there exist nonsingular $C, D \in M(n, \mathbb{F}')$, where \mathbb{F}' is some extension field of \mathbb{F} , such that for all $B \in \mathcal{B}$, the matrix DBC^{-1} is upper-triangular.

To ease the description of our results, we make a few definitions and notations. We denote by $\text{rank}(B)$ the rank of a matrix B , and we set $\text{corank}(B) = n - \text{rank}(B)$. For a matrix space \mathcal{B} we set $\text{rank}(\mathcal{B}) = \max\{\text{rank}(B) \mid B \in \mathcal{B}\}$ and $\text{corank}(\mathcal{B}) = n - \text{rank}(\mathcal{B})$. We say that \mathcal{B} is *singular* if $\text{rank}(\mathcal{B}) < n$, that is if \mathcal{B} does not contain a nonsingular element, and *nonsingular*

¹ At present, we are not aware of the deterministic complexity of computing a rank-1 basis for matrix spaces in \mathbf{R}_1 . Gurvits made a similar comment in [14].

otherwise. For a subspace $U \leq \mathbb{F}^n$, we set $\mathcal{B}(U) = \langle B(u) \mid B \in \mathcal{B}, u \in U \rangle$. Let c be a nonnegative integer. We say that U is a c -singularity witness of \mathcal{B} , if $\dim(U) - \dim(\mathcal{B}(U)) \geq c$, and U is a singularity witness of \mathcal{B} if for some $c > 0$, it is a c -singularity witness.

Note that if there exists a singularity witness of \mathcal{B} then \mathcal{B} can only be singular. Let us define the *discrepancy* of \mathcal{B} as $\text{disc}(\mathcal{B}) = \max\{c \in \mathbb{N} \mid \exists c\text{-singularity witness of } \mathcal{B}\}$. Then it is also clear that $\text{corank}(\mathcal{B}) \geq \text{disc}(\mathcal{B})$. We now state our main theorems.

► **Theorem 1.** *Let \mathbb{F} be either \mathbb{Q} or a finite field. There is a deterministic polynomial-time algorithm which solves the SMR if \mathcal{B} is spanned by rank-1 matrices. If the size of the field \mathbb{F} is at least $n + 1$, the algorithm solves the constructive SMR, and it also outputs a $\text{corank}(\mathcal{B})$ -singularity witness.*

► **Theorem 2.** *Let \mathbb{F} be either \mathbb{Q} or a finite field of size at least $n + 1$. There is a deterministic polynomial-time algorithm which solves the constructive SDIT if \mathcal{B} is triangularizable. Furthermore, over finite fields, when \mathcal{B} is singular it also outputs a singularity witness.*

We remark that Theorem 1 remains true if we weaken the assumptions by only requiring that \mathcal{B} is rank-1 spanned over some extension field of \mathbb{F} rather than over \mathbb{F} . Also, instead of assuming that the whole space \mathcal{B} is rank-1 spanned it is sufficient to suppose that a subspace of \mathcal{B} of co-dimension one is spanned by rank-1 matrices. While the first extension can be achieved easily, the second extension requires some more work (though mostly technical).

1.1 Comparison with previous works

The idea of singularity witnesses was already present in Lovász's work [22]. Among other things, Lovász showed that for the rank-1 spanned case, the equality $\text{corank}(\mathcal{B}) = \text{disc}(\mathcal{B})$ holds, by reducing it to Edmonds' Matroid Intersection theorem [9], which in turn can be deduced from Rado's matroidal generalization of Hall's theorem [24] (see also [28]). Inspired by this fact, Gurvits defined the *Edmonds-Rado property* as the class of matrix spaces which are either nonsingular, or have a singularity witness. He listed several subclasses of the Edmonds-Rado class, including \mathbf{R}_1 (by the aforementioned result of Lovász) and triangularizable matrices. A well-known example of a matrix space without the Edmonds-Rado property is the linear space of skew symmetric matrices of size 3 [22].

As we stated already, Gurvits has presented a polynomial-time deterministic algorithm for the SDIT over \mathbb{Q} for matrix spaces with the Edmonds-Rado property. Therefore over \mathbb{Q} , his algorithm covers the SDIT for \mathbf{R}_1 and for triangularizable matrices. Our algorithms are valid not only over \mathbb{Q} but also over finite fields. In the triangularizable case we also deal with the SDIT, but for \mathbf{R}_1 we solve the more general SMR. In fact, it is not hard to reduce SMR for the general to SMR for the triangularizable case (see Lemma 26 in [18]), so solving SMR for the triangularizable case is as hard as the general case. In both cases the algorithms solve the constructive version of the problems, and they also construct singularity witnesses, except for the SDIT over the rationals. Finally, they work in polynomial time when the field size is at least $n + 1$. Moreover, for \mathbf{R}_1 the algorithm solves the non constructive SMR in polynomial time regardless of the field size, settling the open problem of Gurvits.

Over fields of constant size, the SMR has certain practical implications [16, 17], but is shown to be NP-hard [5] in general. Some special cases have been studied, mostly in the form of the *mixed matrices*, that is linear matrices where each entry is either a variable or a field element. Then by restricting the way variables appear in the matrices some cases turn out to have efficient deterministic algorithms, including when every variable appears at most once ([16], building on [12, 23]), and when the mixed matrix is skew-symmetric

and every variable appears at most twice ([13, 11]). Finally in [19], Ivanyos, Karpinski and Saxena present a deterministic polynomial-time algorithm for the case when among the input matrices B_1, \dots, B_m all but B_1 are of rank 1.

As a computational model of polynomials, determinants with affine polynomial entries turn out to be equivalent to algebraic branching programs (ABPs) [27, 4] up to a polynomial overhead. Thus the identity test for ABPs is the same as SDIT. For restricted classes of ABPs, (quasi)polynomial-time deterministic identity test algorithms have been devised (cf. [10] and the references therein). Note that identity test results for SDIT and ABPs are in general incomparable. For an application of SDIT to quantum information processing see [6].

Let us comment briefly on the main technical tool we use in our algorithms. We generalize the first and second Wong sequences for matrix pencils (essentially two-dimensional matrix spaces) which have turned out to be useful among others in the area of linear differential-algebraic equations (see the recent survey [26]). These were originally defined in [29] for a pair of matrices (A, B) , and were recently used to compute the Kronecker normal form in a numerical stable way [2, 3]. We generalize Wong sequences to the case $(\mathcal{A}, \mathcal{B})$ where \mathcal{A} and \mathcal{B} are matrix spaces, and show that they have analogous basic properties to the original ones. We relate the generalized Wong sequences to Edmonds' problems via singularity witnesses. Essentially this connection allows us to design the algorithm for \mathbf{R}_1 using the second Wong sequence, and the algorithm for triangularizable matrix spaces using the first Wong sequence. We remark that techniques similar to the second Wong sequence were already used in [19].

Organization. In Section 2 we define Wong sequences of a pair of matrix spaces, and present their basic properties. In Section 3 the connection between the second Wong sequence and singularity witnesses is shown. Based on this connection we introduce the power overflow problem, and reduce the SMR to it. We also prove here Theorem 1 under the hypothesis that there is a polynomial time algorithm for the power overflow problem. In Section 4 we show an algorithm for the power overflow problem that works in polynomial time for rank-1 spanned matrix spaces. In Section 5 the algorithm for Theorem 2 is outlined, which works for triangularizable matrix spaces. The readers are referred to the full version [18] for certain missing details, and some discussion on the Edmonds-Rado class and some subclasses.

2 Wong sequences for pairs of matrix spaces

For $n \in \mathbb{N}$, we set $[n] = \{1, \dots, n\}$. We use 0 to denote the zero vector space. In this section we generalize the classical Wong sequences of matrix pencils to the situation of pairs of matrix subspaces. This is the main technical tool in this work. Let V and V' be finite dimensional vector spaces over a field \mathbb{F} , and let $\text{Lin}(V, V')$ be the vector space of linear maps from V to V' . We set $n = \dim(V)$ and $n' = \dim(V')$. For $A \in \text{Lin}(V, V')$, and linear subspaces $\mathcal{A} \leq \text{Lin}(V, V')$, $U \leq V$ and $W \leq V'$, we define $A(U) = \{A(u) \mid u \in U\}$, $\mathcal{A}(U) = \langle \{A(u) \mid A \in \mathcal{A}, u \in U\} \rangle$, $A^{-1}(W) = \{v \in V \mid A(v) \in W\}$, and $\mathcal{A}^{-1}(W) = \{v \in V \mid \forall A \in \mathcal{A}, A(v) \in W\}$. Observe that $A(U)$, $\mathcal{A}(U)$ are linear subspaces of V' , whereas $A^{-1}(W)$ and $\mathcal{A}^{-1}(W)$ are subspaces of V . Also note that $\mathcal{A}(U) = \langle \cup_{A \in \mathcal{A}} A(U) \rangle$ and $\mathcal{A}^{-1}(W) = \cap_{A \in \mathcal{A}} A^{-1}(W)$. Moreover, if \mathcal{A} is spanned by $\{A_1, \dots, A_m\}$, then $\mathcal{A}(U) = \langle \cup_{i \in [m]} A_i(U) \rangle$, and $\mathcal{A}^{-1}(W) = \cap_{i \in [m]} A_i^{-1}(W)$. Some easy and useful facts are the following.

► **Fact 3.** For $\mathcal{A}, \mathcal{B} \leq \text{Lin}(V, V')$, and $U, S \leq V$, $W, T \leq V'$, we have:

1. If $U \subseteq S$ and $W \subseteq T$, then $\mathcal{A}(U) \subseteq \mathcal{A}(S)$ and $\mathcal{A}^{-1}(W) \subseteq \mathcal{A}^{-1}(T)$;
2. If $\mathcal{B}(U) \subseteq \mathcal{A}(U)$ and $\mathcal{B}(S) \subseteq \mathcal{A}(S)$, then $\mathcal{B}(\langle U \cup S \rangle) \subseteq \mathcal{A}(\langle U \cup S \rangle)$;
3. If $\mathcal{B}^{-1}(W) \supseteq \mathcal{A}^{-1}(W)$ and $\mathcal{B}^{-1}(T) \supseteq \mathcal{A}^{-1}(T)$, then $\mathcal{B}^{-1}(W \cap T) \supseteq \mathcal{A}^{-1}(W \cap T)$;
4. $\mathcal{A}^{-1}(\mathcal{A}(U)) \supseteq U$, and $\mathcal{A}(\mathcal{A}^{-1}(W)) \subseteq W$.

We now define the two Wong sequences for a pair of matrix subspaces.

► **Definition 4.** Let $\mathcal{A}, \mathcal{B} \leq \text{Lin}(V, V')$. The sequence of subspaces $(U_i)_{i \in \mathbb{N}}$ of V is called the *first Wong sequence of $(\mathcal{A}, \mathcal{B})$* , where $U_0 = V$, and $U_{i+1} = \mathcal{B}^{-1}(\mathcal{A}(U_i))$. The sequence of subspaces $(W_i)_{i \in \mathbb{N}}$ of V' is called the *second Wong sequences of $(\mathcal{A}, \mathcal{B})$* , where $W_0 = 0$, and $W_{i+1} = \mathcal{B}(\mathcal{A}^{-1}(W_i))$.

When $\mathcal{A} = \langle A \rangle$ and $\mathcal{B} = \langle B \rangle$ are one dimensional matrix spaces, the Wong sequences for $(\mathcal{A}, \mathcal{B})$ coincide with the classical Wong sequences for the matrix pencil $Ax - B$ [29, 2]. The following properties are straightforward generalizations of those for classical Wong sequences. We start by considering the first Wong sequence.

► **Proposition 5.** Let $(U_i)_{i \in \mathbb{N}}$ be the first Wong sequence of $(\mathcal{A}, \mathcal{B})$. Then for all $i \in \mathbb{N}$, we have $U_{i+1} \subseteq U_i$. Furthermore, $U_{i+1} = U_i$ if and only if $\mathcal{B}(U_i) \subseteq \mathcal{A}(U_i)$.

Proof. Firstly we show that $U_{i+1} \subseteq U_i$, for every $i \in \mathbb{N}$. For $i = 0$, this holds trivially. For $i > 0$, by Fact 3 (1) we get $U_{i+1} = \mathcal{B}^{-1}(\mathcal{A}(U_i)) \subseteq \mathcal{B}^{-1}(\mathcal{A}(U_{i-1})) = U_i$, since $U_i \subseteq U_{i-1}$.

Suppose now that $\mathcal{B}(U_i) \subseteq \mathcal{A}(U_i)$, for some i . Then $U_i \subseteq \mathcal{B}^{-1}(\mathcal{B}(U_i)) \subseteq \mathcal{B}^{-1}(\mathcal{A}(U_i))$ respectively by Fact 3 (4) and (1), which gives $U_{i+1} = U_i$. If $\mathcal{B}(U_i) \not\subseteq \mathcal{A}(U_i)$ then there exist $B \in \mathcal{B}$ and $v \in U_i$ such that $B(v) \notin \mathcal{A}(U_i)$. Thus $v \notin \mathcal{B}^{-1}(\mathcal{A}(U_i)) = U_{i+1}$, which gives $U_{i+1} \subset U_i$. ◀

Given Proposition 5, we see that the first Wong sequence stabilizes after at most n steps at some subspace. That is, for any $(\mathcal{A}, \mathcal{B})$, there exists $\ell \in \{0, \dots, n\}$, such that $U_0 \supset U_1 \supset \dots \supset U_\ell = U_{\ell+1} = \dots$. In this case we call the subspace U_ℓ the *limit* of $(U_i)_{i \in \mathbb{N}}$, and we denote it by U^* .

► **Proposition 6.** U^* is the largest subspace $T \leq V$ such that $\mathcal{B}(T) \subseteq \mathcal{A}(T)$.

Proof. By Proposition 5 we know that U^* satisfies $\mathcal{B}(U^*) \subseteq \mathcal{A}(U^*)$. Consider an arbitrary $T \leq V$ such that $\mathcal{B}(T) \subseteq \mathcal{A}(T)$, we show by induction that $T \subseteq U_i$, for all i . When $i = 0$ this trivially holds. Suppose that $T \subseteq U_i$, for some i . Then by repeated applications of Fact 3 we have $T \subseteq \mathcal{B}^{-1}(\mathcal{B}(T)) \subseteq \mathcal{B}^{-1}(\mathcal{A}(T)) \subseteq \mathcal{B}^{-1}(\mathcal{A}(U_i)) = U_{i+1}$. ◀

Analogous properties hold for the second Wong sequence $(W_i)_{i \in \mathbb{N}}$. In particular the sequence stabilizes after at most n' steps, and there exists a *limit* subspace W^* of $(W_i)_{i \in \mathbb{N}}$. We summarize them in the following proposition.

► **Proposition 7.** Let $(W_i)_{i \in \mathbb{N}}$ be the second Wong sequence of $(\mathcal{A}, \mathcal{B})$. Then

1. $W_{i+1} \supseteq W_i$, for all $i \in \mathbb{N}$. Furthermore, $W_{i+1} = W_i$ if and only if $\mathcal{B}^{-1}(W_i) \supseteq \mathcal{A}^{-1}(W_i)$.
2. The limit subspace W^* is the smallest subspace $T \leq V'$ s.t. $\mathcal{B}^{-1}(T) \supseteq \mathcal{A}^{-1}(T)$.

It is worth noting that the second Wong sequence can be viewed as the dual of the first one in the following sense. Assume that V and V' are equipped with nonsingular symmetric bilinear forms, both denoted by $\langle \cdot, \cdot \rangle$. For a linear map $A : V \rightarrow V'$ let $A^T : V' \rightarrow V$ stand for the transpose of A with respect to $\langle \cdot, \cdot \rangle$. This is the unique map with the property $\langle A^T(u), v \rangle = \langle u, A(v) \rangle$, for all $u \in V'$ and $v \in V$. For a matrix space \mathcal{A} , let \mathcal{A}^T be the space $\{A^T | A \in \mathcal{A}\}$. For $U \leq V$, the orthogonal subspace of U is defined as $U^\perp = \{v \in V \mid \langle v, u \rangle = 0 \text{ for all } u \in U\}$. Similarly we define W^\perp for $W \leq V'$. Then we have $((\mathcal{A}^T)^{-1}(U))^\perp = \mathcal{A}(U^\perp)$, and $(\mathcal{A}^T(V))^\perp = \mathcal{A}^{-1}(V^\perp)$. It can be verified that if $(W_i)_{i \in \mathbb{N}}$ is the second Wong sequence of $(\mathcal{A}, \mathcal{B})$ and $(U_i)_{i \in \mathbb{N}}$ the first Wong sequence of $(\mathcal{A}^T, \mathcal{B}^T)$, then $W_i = U_i^\perp$. We note that the duality of Wong sequences was, already derived in [2] for pairs of matrices.

For a matrix space \mathcal{A} and a subspace $U \leq V$ given in terms of a basis we can compute $\mathcal{A}(U)$ by applying the basis elements for \mathcal{A} to those of U and then selecting a maximal set of

linearly independent vectors. A possible way of computing $\mathcal{A}^{-1}(U)$ for $U \leq V'$ is to compute first U^\perp , then $\mathcal{A}^T(U^\perp)$ and finally $\mathcal{A}^{-1}(U) = (\mathcal{A}^T(U^\perp))^\perp$. Therefore we have

► **Proposition 8.** Wong sequences can be computed using $(n + n')^{O(1)}$ arithmetic operations.

Unfortunately, we are unable to prove that over the rationals the bit length of the entries of the bases describing the Wong sequences remain polynomially bounded in the length of the data for \mathcal{A} and \mathcal{B} . However, in Section 3.1 we show that if $\mathcal{A} = \langle A \rangle$, then the first few members of the second Wong sequence which happen to be contained in $\text{im}(A)$ can be computed in polynomial time using an iteration of multiplying vectors by matrices from a basis for \mathcal{B} and by a pseudo-inverse of A .

We also observe that if we consider the bases for \mathcal{A} and \mathcal{B} as matrices over an extension field \mathbb{F}' of \mathbb{F} then the members of the Wong sequences over \mathbb{F}' are just the \mathbb{F}' -linear spaces spanned by the corresponding members of the Wong sequences over \mathbb{F} . In particular, the limit of the first Wong sequence over \mathbb{F} is nontrivial if and only if the limit of the first Wong sequence over \mathbb{F}' is nontrivial.

3 The second Wong sequence and singularity witnesses

3.1 The connection

As in Section 2, let V and V' be finite dimensional vector spaces over a field \mathbb{F} , of respective dimensions n and n' . For $A \in \text{Lin}(V, V')$ we set $\text{corank}(A) = \dim(\ker(A))$. For $\mathcal{B} \leq \text{Lin}(V, V')$, the concepts of c -singularity witnesses, $\text{disc}(\mathcal{B})$ and $\text{corank}(\mathcal{B})$, defined for the case when $n = n'$, can be generalized naturally to \mathcal{B} . We also have that $\text{corank}(\mathcal{B}) \geq \text{disc}(\mathcal{B})$, and that a $\text{corank}(\mathcal{B})$ -singularity witness of \mathcal{B} does not exist necessarily. Let $A \in \mathcal{B}$, and consider $(W_i)_{i \in \mathbb{N}}$, the second Wong sequence of (A, \mathcal{B}) . The next lemma states that the limit W^* is basically such a witness under the condition that it is contained in the image of A . Moreover, in this specific case the limit can be computed efficiently.

► **Lemma 9.** *Let $A \in \mathcal{B} \leq \text{Lin}(V, V')$, and let W^* be the limit of the second Wong sequence of (A, \mathcal{B}) . There exists a $\text{corank}(A)$ -singularity witness of \mathcal{B} if and only if $W^* \subseteq \text{im}(A)$. If this is the case, then A is of maximum rank and $A^{-1}(W^*)$ is a $\text{corank}(\mathcal{B})$ -singularity witness.*

Proof. We prove the equivalence. Firstly suppose that $W^* \subseteq \text{im}(A)$. Then $\dim(A^{-1}(W^*)) = \dim(W^*) + \dim(\ker(A))$. Since $W^* = \mathcal{B}(A^{-1}(W^*))$ and $\dim(\ker(A)) = \text{corank}(A)$, it follows that $A^{-1}(W^*)$ is a $\text{corank}(A)$ -singularity witness of \mathcal{B} .

Let us now suppose that some $U \leq V$ is a $\text{corank}(A)$ -singularity witness, that is $\dim(U) - \dim(\mathcal{B}(U)) \geq \text{corank}(A)$. Then $\dim(U) - \dim(A(U)) \geq \text{corank}(A)$ because $A \in \mathcal{B}$. Since the reverse inequality always holds without any condition on U , we have $\dim(U) - \dim(A(U)) = \text{corank}(A)$. Similarly we have $\dim(U) - \dim(\mathcal{B}(U)) = \text{corank}(A)$ which implies that $\dim(A(U)) = \dim(\mathcal{B}(U))$, and therefore $A(U) = \mathcal{B}(U)$. For a subspace $S \leq V$ the equality $\dim(S) - \dim(A(S)) = \text{corank}(S)$ is equivalent to $\ker(A) \subseteq S$, thus we have $\ker(A) \subseteq U$ from which it follows that $U = A^{-1}(A(U))$. But then $\mathcal{B}^{-1}(A(U)) = \mathcal{B}^{-1}(\mathcal{B}(U)) \supseteq U = A^{-1}(A(U))$. Since W^* is the smallest subspace $T \leq V'$ satisfying $\mathcal{B}^{-1}(T) \supseteq A^{-1}(T)$, we can conclude that $W^* \subseteq A(U)$.

The existence of a $\text{corank}(A)$ -singularity witness obviously implies that A is of maximum rank, and when $W^* \subseteq \text{im}(A)$ we have already seen that $A^{-1}(W^*)$ is a $\text{corank}(A)$ -singularity witness of \mathcal{B} . Since $\text{corank}(A) = \text{corank}(\mathcal{B})$, it is also a $\text{corank}(\mathcal{B})$ -singularity witness. ◀

We would like to find an efficient way of testing whether $W^* \subseteq \text{im}(A)$ for a given $A \in \mathcal{B}$. In the computation of the limit W^* of the second Wong sequence of (A, \mathcal{B}) the

computationally hard step is applying iteratively A^{-1} . We overcome this difficulty by introducing a pseudo-inverse of A in the computation. We describe now this method.

Let $n = \dim(V)$ and $n' = \dim(V')$. First of all we assume without loss of generality that $n = n'$. Indeed, if $n < n'$ we can add as a direct complement a suitable space to V on which \mathcal{B} acts as zero, and if $n > n'$, we can embed V' into a larger space. In terms of matrices, this means augmenting the elements of \mathcal{B} by zero columns or zero rows to obtain square matrices. This procedure affects neither the ranks of the matrices in \mathcal{B} nor the singularity witnesses.

We say that a nonsingular linear map $A' : V' \rightarrow V$ is a *pseudo-inverse* of A if the restriction of A' to $\text{im}(A)$ is the inverse of the restriction of A to a direct complement of $\ker(A)$. Such a map can be efficiently constructed as follows. Choose a direct complement U of $\ker(A)$ in V as well as a direct complement U' of $\text{im}(A)$ in V' . Then take the map $A'_0 : \text{im}(A) \rightarrow U$ such that AA'_0 is the identity of $\text{im}(A)$ and take an arbitrary nonsingular linear map $A'_1 : U' \rightarrow \ker(A)$. Finally let A' be the direct sum of A'_0 and A'_1 .

► **Lemma 10.** *Let $A \in \mathcal{B} \leq \text{Lin}(V, V')$ and let A' be a pseudo-inverse of A . There exists a $\text{corank}(A)$ -singularity witness of \mathcal{B} if and only if $(\mathcal{B}A')^i(\ker(AA')) \subseteq \text{im}(A)$, for all $i \in [n]$. This can be tested in polynomial time, and if the condition holds then A is of maximum rank and $A'(W^*)$ is a $\text{corank}(\mathcal{B})$ -singularity witness which also can be computed deterministically in polynomial time.*

Proof. It follows from Lemma 9 that a $\text{corank}(A)$ -singularity witness exists if and only if $W_i \subseteq \text{im}(A)$, for $i = 1, \dots, n$. Observing that $(\mathcal{B}A')^i(\ker(AA')) \subseteq W_i$ for $i = 1, \dots, n$, to prove the equivalence it is sufficient to show that if $(\mathcal{B}A')^i(\ker(AA')) \subseteq \text{im}(A)$ for $i = 1, \dots, n$ then $W_i = (\mathcal{B}A')^i(\ker(AA'))$ for $i = 1, \dots, n$. The proof is by induction. For $i = 1$ the claim $W_1 = \mathcal{B}A'(\ker(AA'))$ holds since $\ker(AA') = A'^{-1}(\ker(A))$. For $i > 1$, by definition $W_i = \mathcal{B}A^{-1}(W_{i-1})$. Since every subspace $W \leq \text{im}(A)$ satisfies $A^{-1}W = A'W + \ker(A)$, where $+$ denotes the direct sum, we get $W_i \subseteq \mathcal{B}A'(W_{i-1}) + \mathcal{B}(\ker(A))$. Observe that $\mathcal{B}(\ker(A)) = W_1$. We will show that $W_1 \subseteq \mathcal{B}A'(W_{i-1})$ and then we conclude by the inductive hypothesis. We know that $W_1 \subseteq W_{i-1}$ from the properties of the Wong sequence, therefore it is sufficient to show that $W_{i-1} \subseteq \mathcal{B}A'(W_{i-1})$. But $W_{i-1} = AA'(W_{i-1})$ since $W_i \subseteq \text{im}(A)$ and A' is the inverse of A on $\text{im}(A)$.

Based on this equivalence, testing the existence of a $\text{corank}(A)$ -singularity witness can be accomplished by a simple algorithm, [18, Lemma 10] for details.

If we find that the condition holds then $A'(W^*)$ by Lemma 9 is a $\text{corank}(\mathcal{B})$ -singularity witness, and it can be easily computed from W^* . ◀

3.2 The power overflow problem

For $A \in \mathcal{B} \leq \text{Lin}(V, V')$, we would like to know whether A is of maximum rank in \mathcal{B} . With the help of the limit W^* of the second Wong sequence of (A, \mathcal{B}) we have established a sufficient condition: we know that if $W^* \subseteq \text{im}(A)$ then A is indeed of maximum rank. Our results until now do not give a necessary condition for the maximum rank. Now we show that the second Wong sequence actually allows to translate this question to the *power overflow* problem (PO) which we define below. As a consequence an efficient solution of the PO guarantees an efficient solution for the SMR. The reduction is mainly based on a theorem of Atkinson and Stephens [1] which essentially says that over big enough fields, in 2-dimensional matrix spaces \mathcal{B} , the equality $\text{corank}(\mathcal{B}) = \text{disc}(\mathcal{B})$ holds.

► **Proposition 11 ([1]).** Assume that $|\mathbb{F}| > n$, and let $A, B \in \text{Lin}(V, V')$. If A is a maximum rank element of $\langle A, B \rangle$ then there exists a $\text{corank}(A)$ -singularity witness of $\langle A, B \rangle$.

Combining Lemma 10 and Proposition 11 we get also an equivalent condition for A being of maximum rank.

► **Lemma 12.** *Assume that $|\mathbb{F}| > n$. Let $A \in \mathcal{B} \leq \text{Lin}(V, V')$, and let A' be a pseudo-inverse of A . Then A is of maximum rank in \mathcal{B} if and only if for every $B \in \mathcal{B}$ and for all $i \in [n]$, we have*

$$(BA')^i(\ker(AA')) \subseteq \text{im}(A).$$

Proof. First observe that A is of maximum rank in \mathcal{B} if and only if for every $B \in \mathcal{B}$, it is of maximum rank in $\langle A, B \rangle$. For a fixed B , by Proposition 11 and Lemma 10, A is of maximum rank in $\langle A, B \rangle$ exactly when $(\langle B, A \rangle A')^i(\ker(AA')) \subseteq \text{im}(A)$, for all $i \in [n]$. From that we can conclude since A' is the inverse of A on $\text{im}(A)$. ◀

This lemma leads us to reduce the problems of deciding if A is of the maximum rank, and finding a matrix of rank larger than A when this is not the case, to the following question.

► **Problem 13 (The power overflow problem).** Given $\mathcal{D} \leq M(n, \mathbb{F})$, $U \leq \mathbb{F}^n$ and $U' \leq \mathbb{F}^n$, output $D \in \mathcal{D}$ and $\ell \in [n]$ s.t. $D^\ell(U) \not\subseteq U'$, if there exists such (D, ℓ) . Otherwise say **no**.

The power overflow problem admits an efficient randomized algorithm when $|\mathbb{F}| = \Omega(n)$. For the rank-1 spanned case we show a deterministic solution regardless of the field size.

► **Theorem 14.** *Let $\mathcal{D} \leq M(n, \mathbb{F})$ be spanned by rank-1 matrices. Then there exists $D \in \mathcal{D}$ and $\ell \in [n]$ such that $D^\ell(U) \not\subseteq U'$ if and only if there exists $\ell \in [n]$ such that $\mathcal{D}^\ell(U) \not\subseteq U'$. The power overflow problem for \mathcal{D} can be solved deterministically in polynomial time.*

Using this result whose proof is given in Section 4 we are now ready to prove Theorem 1.

Proof of Theorem 1. First we suppose that $|\mathbb{F}| \geq n + 1$. Let A be an arbitrary matrix in \mathcal{B} . The algorithm iterates the following process until A becomes of maximum rank.

We run the algorithm of Lemma 10 to test whether $(\mathcal{B}A')^i(\ker(AA')) \subseteq \text{im}(A)$ for $i \in [n]$. If this condition holds then A is of maximum rank, and the algorithm also gives a corank(\mathcal{B})-singularity witness. Otherwise we know by Theorem 14 that there exists $B \in \mathcal{B}$ and $i \in [n]$ such that $(BA')^i(\ker(AA')) \not\subseteq \text{im}(A)$. We apply the algorithm of Theorem 14 with input $\mathcal{B}A'$, $\ker(AA')$ and $\text{im}(A)$, which finds such a couple (B, i) . Lemma 12 applied to $\langle A, B \rangle$ implies that A is not of maximum rank in $\langle A, B \rangle$. If A has rank $r \leq n - 1$ which is not maximal in $\langle A, B \rangle$, then the determinant of an appropriate $(r + 1) \times (r + 1)$ minor is a nonzero polynomial of degree at most $r + 1$ which has at most $r + 1 \leq n$ roots. We then pick $n + 1$ arbitrary field elements $\lambda_1, \dots, \lambda_{n+1}$, and we know that for some $1 \leq j \leq n + 1$ we have $\text{rank}(A + \lambda_j B) > \text{rank}(A)$. We replace A by $A + \lambda_j B$ and restart the process.

At the end of each iteration, by a reduction procedure described in [7] we can achieve that the matrix A , written as a linear combination of B_1, \dots, B_m has coefficients from a fixed subset $K \subseteq \mathbb{F}$ of size $n + 1$. In fact, if $A = \alpha_1 B_1 + \alpha_2 B_2 \dots + \alpha_m B_m$ has rank r then for at least one $\kappa_1 \in K$ the matrix $\kappa_1 B_1 + \alpha_2 B_2 \dots + \alpha_m B_m$ has rank at least r . This way all the coefficients α_j can be replaced with an appropriate element from K .

As in each iteration we either stop (and conclude with A being of maximum rank), or increase the rank of A by at least 1, the number of iterations is at most n . Also, each iteration takes polynomial many steps since the processes of Lemma 10 and Theorem 14 are polynomial. Therefore the overall running time is also polynomial. ◀

We can compute the maximum rank over a field of size less than $n + 1$ by running the above procedure over a sufficiently large extension field. The maximum rank will not grow if we go over an extension. This follows from the fact that the equality $\text{corank}(\mathcal{B}) = \text{disc}(\mathcal{B})$ holds over any field if \mathcal{B} is spanned by an arbitrary matrix and by rank one matrices, see [19].

4 The power overflow problem for rank-1 spanned matrix spaces

In this section we prove **Theorem 14**. Given subspaces U, U' of \mathbb{F}^n as well as a basis $\{D_1, \dots, D_m\}$ for a matrix space $\mathcal{D} \leq M(n, \mathbb{F})$, we will show is that in polynomial time we can decide if $\mathcal{D}^\ell(U) \not\subseteq U'$ for some ℓ , and if this holds then find $D \in \mathcal{D}$ s.t. $D^\ell(U) \not\subseteq U'$.

Formally let $\ell = \ell(\mathcal{D})$ be the smallest integer j s.t. $\mathcal{D}^j(U) \not\subseteq U'$ if such an integer exists, and n otherwise. We start by computing ℓ and for $1 \leq j \leq \ell$, bases \mathcal{T}_j for \mathcal{D}^j . Set $\mathcal{T}_1 = \{D_1, \dots, D_m\}$. If $\mathcal{D}^j(U) \not\subseteq U'$ then we set $\ell = j$ and stop constructing further bases. If $j = n$ and $\mathcal{D}^n(U) \subseteq U'$ then we stop the algorithm and output **no**. Otherwise we compute \mathcal{T}_{j+1} by selecting a maximal linearly independent set from the products of elements in \mathcal{T}_j and \mathcal{T}_1 .

We are now looking for D such that $D^\ell(U) \not\subseteq U'$. For $i \in [\ell]$, we define subspaces \mathcal{H}_i of \mathcal{D} , which play a crucial role in the algorithm:

$$\mathcal{H}_i = \{X \in \mathcal{D} \mid \mathcal{D}^{\ell-j} X \mathcal{D}^{j-1}(U) \subseteq U', j = 1, \dots, i-1, i+1, \dots, \ell\}.$$

That is, $X \in \mathcal{H}_i$ if and only if whenever X appears in a place other than the i th in a product P of ℓ elements from \mathcal{D} then $P(U) \subseteq U'$. The subspaces \mathcal{H}_i can be computed as follows. Let x_1, \dots, x_m be formal variables, an element in \mathcal{D} can be written as $X = \sum_{k \in [m]} x_k D_k$. The condition $\mathcal{D}^{\ell-j} X \mathcal{D}^{j-1}(U) \subseteq U'$ is equivalent to the set of the following homogeneous linear equations in the variables x_k : $\langle Z(\sum_{k \in [m]} x_k D_k) Z' u, v \rangle = 0$, where Z is from $\mathcal{T}_{\ell-j}$, Z' is from \mathcal{T}_{j-1} , u is from a basis for U and v is from a basis for U'^\perp . Thus \mathcal{H}_i can be computed by solving a system of polynomially many homogeneous linear equations. Note that the coefficients of the equations are scalar products of vectors from a basis for U'^\perp by vectors obtained as applying products of ℓ matrices from $\{D_1, \dots, D_m\}$ to basis elements for U . The definition of \mathcal{H}_i implies the following.

► **Lemma 15.** *For a matrix $X = X_1 + \dots + X_\ell$ with $X_i \in \mathcal{H}_i$, we have $X^\ell(U) \subseteq U'$ if and only if $X_\ell \cdots X_2 X_1(U) \subseteq U'$.*

Proof. We have $X^m = \sum_{\sigma} X_{\sigma(\ell)} \cdots X_{\sigma(1)}$, where the summation is over the maps $\sigma : [\ell] \rightarrow [\ell]$. When σ is not the identity map then there exists an index j such that $\sigma(j) \neq j$. Then $X_{\sigma(\ell)} \cdots X_{\sigma(1)}(U) \subseteq U'$ by the definition of $\mathcal{H}_{\sigma(j)}$. ◀

In general, \mathcal{H}_i can be 0. In our setting, due to the existence of rank one generators, fortunately this is far from the case. Recall that ℓ is the smallest integer such that $\mathcal{D}^\ell(U) \not\subseteq U'$.

► **Lemma 16.** *We have $\mathcal{H}_\ell \cdots \mathcal{H}_1(U) \not\subseteq U'$.*

Proof. Assume that \mathcal{D} is spanned by the rank one matrices C_1, \dots, C_m . Then there exist indices k_1, \dots, k_ℓ such $C_{k_\ell} \cdots C_{k_1}(U) \not\subseteq U'$. We show that $C_{k_i} \in \mathcal{H}_i$, for $i \in [\ell]$, this implies immediately $\mathcal{H}_\ell \cdots \mathcal{H}_1(U) \not\subseteq U'$. Assume by contradiction that $C_{k_i} \notin \mathcal{H}_i$, for some $i \in [\ell]$. Then $\mathcal{D}^{\ell-j} C_{k_i} \mathcal{D}^{j-1}(U) \not\subseteq U'$, for some $j \neq i$. On the other hand C_{k_i} satisfies $\mathcal{D}^{\ell-i} C_{k_i} \mathcal{D}^{i-1}(U) \not\subseteq U'$. Since C_{k_i} is of rank 1 we have $C_{k_i} \mathcal{D}^{j-1}(U) = C_{k_i} \mathcal{D}^{i-1}(U)$, which yields that neither $\mathcal{D}^{\ell-i} C_{k_i} \mathcal{D}^{j-1}(U)$ nor $\mathcal{D}^{\ell-j} C_{k_i} \mathcal{D}^{i-1}(U)$ is contained in U' . However one of these products is shorter than ℓ , contradicting the minimality of ℓ . ◀

To finish the algorithm, we compute bases for products $\mathcal{H}_i \cdots \mathcal{H}_1$, for $i \in [n]$, in a way similar to computing bases for \mathcal{D}^i . Then we search the basis of \mathcal{H}_ℓ for an element Z such that $Z \mathcal{H}_{\ell-1} \cdots \mathcal{H}_1(U) \not\subseteq U'$. We put $X_\ell = Z$ and continue searching the basis of $\mathcal{H}_{\ell-1}$ for an element Z such that $X_\ell Z \mathcal{H}_{\ell-2} \cdots \mathcal{H}_1(U) \not\subseteq U'$. Continuing the iteration, Lemma 16 ensures that eventually we find $X_i \in \mathcal{H}_i$, for $i \in [\ell]$, such that $X_\ell \cdots X_1(U) \not\subseteq U'$. We set $D = X_1 + \dots + X_\ell$, then by Lemma 15 we have $D^\ell(U) \not\subseteq U'$. We return D and ℓ . ◻

5 The first Wong sequence and triangularizable matrix spaces

Here we only give a proof outline of Theorem 2, and the reader is referred to the full version [18, Section 5] for details. Our task is to determine whether there exists a nonsingular matrix in a triangularizable matrix space, and finding such a matrix if exists. Let \mathbb{F}' be an extension field of \mathbb{F} , and recall that $\mathcal{B} \leq M(n, \mathbb{F})$ is triangularizable if there exist nonsingular $C, D \in M(n, \mathbb{F}')$, s.t. $\forall B \in \mathcal{B}$, DBC^{-1} is upper triangular. Our starting point is the following lemma, which connects first Wong sequences with singularity witnesses.

► **Lemma 17.** *Let $A \in \mathcal{B} \leq M(n, \mathbb{F})$, and let U^* be the limit of the first Wong sequence of (A, \mathcal{B}) . Set $d = \dim(U^*)$. Then either U^* is a singularity witness of \mathcal{B} , or there exist nonsingular matrices $P, Q \in M(n, \mathbb{F})$, such that $\forall B \in \mathcal{B}$, QBP^{-1} is of the form $\begin{bmatrix} X & Y \\ 0 & Z \end{bmatrix}$, where X is of size $d \times d$, and \mathcal{B} is nonsingular in the X -block.*

Lemma 17 suggests a recursive algorithm: take an arbitrary $A \in \mathcal{B}$ and compute U^* , the limit of the first Wong sequence of (A, \mathcal{B}) . If we get a singularity witness, we are done. Otherwise, if $U^* \neq 0$, as the X -block is already nonsingular, we only need to focus on the nonsingularity of Z -block which is of smaller size. To make this idea work, we have to satisfy essentially two conditions. We must find some A such that $U^* \neq 0$, and to allow for recursion the specific property of the matrix space \mathcal{B} we are concerned with has to be inherited by the subspace corresponding to the Z -block. It turns out that in the triangularizable case these two problems can be taken care of by the following Lemma.

► **Lemma 18.** *Let $\mathcal{B} \leq \mathbb{F}$ be given by a basis $\{B_1, \dots, B_m\}$, and suppose that there exist nonsingular matrices $C, D \in M(n, \mathbb{F}')$ such that $B_i = DB'_i C^{-1}$ and $B'_i \in M(n, \mathbb{F}')$ is upper triangular for every $i \in [m]$. Then we have the following.*

1. *Either $\bigcap_{i \in [m]} \ker(B_i) \neq 0$, or there exists $j \in [m]$ and $0 \neq U \leq \mathbb{F}^n$ s.t. $B_j(U) = \mathcal{B}(U)$.*
2. *Suppose there exist $j \in [m]$ and $0 \neq U \leq \mathbb{F}^n$ s.t. $B_j(U) = \mathcal{B}(U)$, and $\dim(U) = \dim(\mathcal{B}_j(U))$. Let $B'_i : \mathbb{F}^n/U \rightarrow \mathbb{F}^n/\mathcal{B}(U)$ be the linear map induced by B_i , for $i \in [m]$. Then $\mathcal{B}^* = \langle B'_1, \dots, B'_m \rangle$ is triangularizable over \mathbb{F}' .*

Proof. 1. Let $\{e_i \mid i \in [n]\}$ be the standard basis of \mathbb{F}^n , and $c_i = C(e_i)$ and $d_i = D(e_i)$ for $i \in [n]$. If $B'_i(1, 1) = 0$ for all $i \in [m]$ then c_1 is in the kernel of every B_i 's. If there exists j such that $B'_j(1, 1) \neq 0$, we set $U' = \langle c_1 \rangle \leq \mathbb{F}^n$. Then it is clear that $\langle d_1 \rangle = B_j(U') = \mathcal{B}(U')$. It follows that the first Wong sequence of (B_j, \mathcal{B}) over \mathbb{F}' has nonzero limit, and therefore the same holds over \mathbb{F} . We can choose for U this limit.

2. First we recall that for a vector space V of dimension n , a complete flag of V is a nested sequence of subspaces $0 = V_0 \subset V_1 \subset \dots \subset V_n = V$. For $\mathcal{A} \leq \text{Lin}(V, V')$ with $\dim(V) = \dim(V') = n$, the matrix space \mathcal{A} is triangularizable if and only if \exists complete flags $0 = V_0 \subset V_1 \subset \dots \subset V_n = V$ and $0 = V'_0 \subset V'_1 \subset \dots \subset V'_n = V'$ s.t. $\mathcal{A}(V_i) \subseteq V'_i$ for $i \in [n]$.

For $U \leq \mathbb{F}^n$, let $\mathbb{F}'U$ be the linear span of U in \mathbb{F}'^n . We think of B_i 's and B'_i 's as linear maps over \mathbb{F}' in a natural way. Let $\ell = \dim(\mathbb{F}'^n/\mathbb{F}'U)$. For $0 \leq i \leq n$ set $S_i = \langle c_1, \dots, c_i \rangle$ and $T_i = \langle d_1, \dots, d_i \rangle$. Obviously $\mathcal{B}(S_i) \subseteq T_i$ for $0 \leq i \leq n$. Let $S_i^* = S_i/\mathbb{F}'U$ and $T_i^* = T_i/\mathcal{B}(\mathbb{F}'U)$, and consider $S_0^* \subseteq \dots \subseteq S_n^*$ and $T_0^* \subseteq \dots \subseteq T_n^*$. We claim that $\forall i \in [n]$, $\dim(S_i^*) \geq \dim(T_i^*)$. This is because as $T_i \cap \mathcal{B}(\mathbb{F}'U) \supseteq B_j(S_i \cap \mathbb{F}'U)$, by $\dim(\mathbb{F}'U) = \dim(\mathcal{B}_j(\mathbb{F}'U))$, $\dim(B_j(S_i \cap \mathbb{F}'U)) \geq \dim(S_i \cap \mathbb{F}'U)$. Thus $\dim(S_i \cap \mathbb{F}'U) \leq \dim(T_i \cap \mathcal{B}(\mathbb{F}'U))$, and $\dim(S_i^*) \geq \dim(T_i^*)$. As $\mathcal{B}^*(S_i^*) \subseteq T_i^*$, $\dim(S_{i+1}^*) - \dim(S_i^*) \leq 1$, and $\dim(T_{i+1}^*) - \dim(T_i^*) \leq 1$, there exist two nested sequences $S_0^* \subset S_{j_1}^* \subset \dots \subset S_{j_\ell}^* = S_n^*$ and $T_0^* \subset T_{k_1}^* \subset \dots \subset T_{k_\ell}^* = T_n^*$, s.t. $\dim(S_{j_h}^*) = \dim(T_{k_h}^*) = h$. Furthermore, by $\dim(S_i^*) \geq \dim(T_i^*)$, $j_h \leq k_h$, thus

$\mathcal{B}^*(S_{j_h}^*) \subseteq \mathcal{B}^*(S_{k_h}^*) \subseteq T_{k_h}^*$, $\forall h \in [\ell]$. That is, the two nested sequences are complete flags, and \mathcal{B}^* is triangularizable over \mathbb{F}' . \blacktriangleleft

Given the above preparation, we can now outline the algorithm for Theorem 2.

Proof of Theorem 2. First we consider finite fields. The algorithm recurses on the size of the matrices, with the base case being the size one. It checks at the beginning whether $\bigcap_{i \in [m]} \ker(B_i) = 0$. If this is the case then it returns $\bigcap_{i \in [m]} \ker(B_i)$ which is a singularity witness. Otherwise, for all $i \in [m]$, it computes the limit U_i^* of the first Wong sequence for (B_i, \mathcal{B}) . By Lemma 18 (1) there exists $j \in [m]$ such that $U_j^* \neq 0$ and $B_j(U_j^*) = \mathcal{B}(U_j^*)$. The algorithm then recurses on the induced actions B_i^* 's of B_i 's, which are also triangularizable by Lemma 18 (2). When \mathcal{B} is nonsingular the algorithm should return a nonsingular matrix. This nonsingular matrix is built step by step by the recursive calls, at each step we have to construct a nonsingular linear combination of B_j and the matrix returned by the call. For this we need $n + 1$ field elements.

The case of the rational numbers can be reduced to the case of finite fields. Let b be a bound on the absolute values of entries in B_i 's. It can be shown that there exists a prime number p of value polynomially bounded by $\log b$ and n s.t. the following holds: let B_i' be the matrix B_i modulo p . When \mathcal{B} is triangularizable and nonsingular then the matrix space spanned by B_i' is triangularizable over an extension field of \mathbb{F}_p and nonsingular. If \mathcal{B} is singular, modulo any prime the matrix space is singular. So we enumerate all prime numbers up to the given polynomial bound, and for each prime use the algorithm over finite fields. \blacktriangleleft

Acknowledgements. We would like to thank the anonymous reviewers for careful reading and pointing out some gaps in an earlier version of the paper. Most of this work was conducted when G. I., Y. Q. and M. S. were at the Centre for Quantum Technologies (CQT) in Singapore, and partially funded by the Singapore Ministry of Education and the National Research Foundation, also through the Tier 3 Grant “Random numbers from quantum processes”. Research partially supported by the European Commission IST STREP project Quantum Algorithms (QALGO) 600700, by the French ANR Blanc program under contract ANR-12-BS02-005 (RDAM project), by the Hungarian Scientific Research Fund (OTKA), Grants NK105645 and K77476, and by the Hausdorff grant EXC59-1/2.

References

- 1 M. D. Atkinson and N. M. Stephens. Spaces of matrices of bounded rank. *The Quarterly Journal of Mathematics*, 29(2):221–223, 1978.
- 2 T. Berger and S. Trenn. The quasi-Kronecker form for matrix pencils. *SIAM Journal on Matrix Analysis and Applications*, 33(2):336–368, 2012.
- 3 Thomas Berger and Stephan Trenn. Addition to “the quasi-Kronecker form for matrix pencils”. *SIAM Journal on Matrix Analysis and Applications*, 34(1):94–101, 2013.
- 4 Stuart J. Berkowitz. On computing the determinant in small parallel time using a small number of processors. *Information Processing Letters*, 18(3):147–150, 1984.
- 5 Jonathan F. Buss, Gudmund S. Frandsen, and Jeffrey O. Shallit. The computational complexity of some problems of linear algebra. *J. Comput. Syst. Sci.*, 58(3):572–596, 1999.
- 6 Eric Chitambar, Runyao Duan, and Yaoyun Shi. Multipartite-to-bipartite entanglement transformations and polynomial identity testing. *Physical Review A*, 81(5):052310, 2010.
- 7 Willem A. de Graaf, Gábor Ivanyos, and Lajos Rónyai. Computing Cartan subalgebras of Lie algebras. *Applicable Algebra in Engineering, Communication and Computing*, 7(5):339–349, 1996.

- 8 Jack Edmonds. Systems of distinct representatives and linear algebra. *J. Res. Nat. Bur. Standards Sect. B*, 71:241–245, 1967.
- 9 Jack Edmonds. Submodular functions, matroids, and certain polyhedra. In N. Sauer R. K. Guy, H. Hanani and J. Schönheim, editors, *Combinatorial Structures and their Appl.*, pages 69–87, New York, 1970. Gordon and Breach.
- 10 Michael A. Forbes and Amir Shpilka. Quasipolynomial-time identity testing of non-commutative and read-once oblivious algebraic branching programs. In *FOCS*, 2013.
- 11 James Geelen and Satoru Iwata. Matroid matching via mixed skew-symmetric matrices. *Combinatorica*, 25(2):187–215, 2005.
- 12 James F. Geelen. Maximum rank matrix completion. *Linear Algebra and its Applications*, 288:211–217, 1999.
- 13 James F. Geelen, Satoru Iwata, and Kazuo Murota. The linear delta-matroid parity problem. *Journal of Combinatorial Theory, Series B*, 88(2):377–398, 2003.
- 14 Leonid Gurvits. Quantum matching theory (with new complexity theoretic, combinatorial and topological insights on the nature of the quantum entanglement), 2002.
- 15 Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. Syst. Sci.*, 69(3):448–484, 2004.
- 16 Nicholas J. A. Harvey, David R. Karger, and Kazuo Murota. Deterministic network coding by matrix completion. In *Proceedings of SODA*, pages 489–498. ACM-SIAM, 2005.
- 17 Nicholas J. A. Harvey, David R. Karger, and Sergey Yekhanin. The complexity of matrix completion. In *Proceedings of SODA*, pages 1103–1111. ACM-SIAM, 2006.
- 18 Gábor Ivanyos, Marek Karpinski, Youming Qiao, and Miklos Santha. Generalized wong sequences and their applications to edmonds' problems. *Electronic Colloquium on Computational Complexity (ECCC)*, 20:103, 2013.
- 19 Gábor Ivanyos, Marek Karpinski, and Nitin Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.
- 20 Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1-2):1–46, 2004.
- 21 László Lovász. On determinants, matchings, and random algorithms. In *FCT*, pages 565–574, 1979.
- 22 László Lovász. Singular spaces of matrices and their application in combinatorics. *Boletim da Sociedade Brasileira de Matemática-Bulletin/Brazilian Mathematical Society*, 20(1):87–99, 1989.
- 23 Kazuo Murota. *Matrices and matroids for systems analysis*. Springer, 2000.
- 24 Richard Rado. A theorem on independence relations. *The Quarterly Journal of Mathematics, Oxford Ser.*, 13(1):83–89, 1942.
- 25 Jacob T. Schwartz. Probabilistic algorithms for verification of polynomial identities. In Edward W. Ng, editor, *Symbolic and Algebraic Computation*, volume 72 of *Lecture Notes in Computer Science*, pages 200–215. Springer Berlin Heidelberg, 1979.
- 26 Stephan Trenn. Solution concepts for linear DAEs: A survey. In Achim Ilchmann and Timo Reis, editors, *Surveys in Differential-Algebraic Equations I*, Differential-Algebraic Equations Forum, pages 137–172. Springer Berlin Heidelberg, 2013.
- 27 Leslie G. Valiant. Completeness classes in algebra. In *STOC*, pages 249–261, 1979.
- 28 D. J. A. Welsh. On matroid theorems of Edmonds and Rado. *Journal of the London Mathematical Society*, 2(2):251–256, 1970.
- 29 Kai-Tak Wong. The eigenvalue problem $\lambda Tx + Sx$. *Journal of Differential Equations*, 16(2):270 – 280, 1974.
- 30 Richard Zippel. Probabilistic algorithms for sparse polynomials. In Edward W. Ng, editor, *Symbolic and Algebraic Computation*, volume 72 of *LNCS*, pages 216–226. Springer, 1979.