Generalized Wong sequences and their applications to Edmonds' problems

Gábor Ivanyos* Marek Karpinski† Youming Qiao‡ Miklos Santha§

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 fields of size at least $n+1$, and the first algorithm actually solves (the non-constructive) SMR independent of the field size. Our framework is based on a generalization of Wong sequences, a classical method to deal with pairs of matrices, to the case of pairs of matrix spaces. [1]

1 Introduction

In 1967, Edmonds introduced the following problem [9]: 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

*Institute for Computer Science and Control, Hungarian Academy of Sciences, Budapest, Hungary (Gabor.Ivanyos@sztaki.mta.hu).

†Department of Computer Science, University of Bonn, Bonn, Germany (marek@cs.uni-bonn.de).

‡Centre for Quantum Computation and Intelligent Systems University of Technology, Sydney; and Centre for Quantum Technologies, National University of Singapore, Singapore 117543. (jimmyqiao86@gmail.com)

§LIAFA, Univ. Paris 7, CNRS, 75205 Paris, France; and Centre for Quantum Technologies, National University of Singapore, Singapore 117543 (miklos.santha@liafa.jussieu.fr).

[1] A preliminary report on this work appeared in [21].

rationals. In this paper we consider this question and its certain variants over more general fields.

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 $\{B_1, \ldots, B_m\} \subseteq M(n, \mathbb{F})$, determine the maximum rank among matrices in $\mathcal{B} = \langle B_1, \ldots, B_m \rangle$, 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 [15] and in [22]). 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 [23] (in [18] 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 [29, 34] it admits an efficient randomized algorithm [25]. Obtaining a deterministic polynomial-time algorithm for the SDIT would be of fundamental importance, since Kabanets and Impagliazzo [23] 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* $B_1, \ldots, B_m$ satisfy certain property. For example, Lovász [26] 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 [15, 27, 19, 16, 14, 22], see Section 1.1 for more details.

Another direction also studied is when instead of the given matrices, the spanned *matrix space* $\mathcal{B} = \langle B_1, \ldots, B_m \rangle$ 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 [18] presented an efficient deterministic algorithm for the SDIT over subfields of $\mathbb{C}$, when the matrix space falls in a special class, what we call the *Edmonds-Rado class*. We shall review the definition of this class in Section 1.1. In this paper, our main goal is to consider Edmonds' problems for the following two classes.

- The class of rank-1 spanned matrix spaces, $\mathbf{R_1}$: a matrix space $\mathcal{B} \leq M(n, \mathbb{F})$ is in $\mathbf{R_1}$, if $\mathcal{B}$ has a basis consisting of rank-1 matrices over $\mathbb{F}'$, where $\mathbb{F}'$ is some extension field of $\mathbb{F}$.[2]

- The class of (upper) triangularizable matrix spaces, $\mathbf{UT}$: a matrix $\mathcal{B} \leq M(n, \mathbb{F})$ is in $\mathbf{UT}$, 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.

It is known that the Edmonds-Rado class includes $\mathbf{R_1}$ and $\mathbf{UT}$. See Section 1.1 for more details. While Gurvits presented an efficient deterministic SDIT algorithm for the Edmonds-Rado class over subfields of $\mathbb{C}$, the same problem over (large enough) finite fields is still open,

[2]Note that it is possible for $\mathcal{B}$ to have a rank-1 basis over $\mathbb{F}'$ but no such over $\mathbb{F}$. See [17] for an example.

even for special classes like $\mathbf{R_1}$ and $\mathbf{UT}$. In fact, Gurvits stated as an open question the complexity of the SMR for $\mathbf{R_1}$ over finite fields [18, page 456].

The difference between properties of matrices and properties of matrix spaces is critical for Edmonds' problems. In particular, whether a matrix space satisfies a certain property or not, should not depend on choices of basis. We are not aware of any result on the complexity of finding rank one generators for a subspace $\mathcal{B}$ in $\mathbf{R_1}$ if it is given by a basis consisting of not necessarily rank one matrices. We believe that the problem is hard. Thus the existence of algorithms for SMR when the B_i's are rank-1 does not immediately imply algorithms for matrix spaces in $\mathbf{R_1}$.

Furthermore, most properties we encounter in practice respect the following equivalence relation of matrix spaces. Two matrix spaces $\mathcal{A}$ and $\mathcal{B}$ in $M(n, \mathbb{F})$ are equivalent, if there exist nonsingular $C, D \in M(n, \mathbb{F})$, s.t. $\mathcal{A} = C\mathcal{B}D := \{CBD \mid B \in \mathcal{B}\}$. Edmonds-Rado class, $\mathbf{R_1}$ and $\mathbf{UT}$ all respect this equivalence relation. Again, given matrices $B_1, \ldots, B_m$, and suppose $\mathcal{B} = \langle B_1, \ldots, B_m \rangle$ is in $\mathbf{UT}$, it is not clear how difficult is computing matrices C, D that triangularize $\mathcal{B}$. The problem does not look as hard as finding rank one generators, see Section 7 for some details. Thus while SDIT for upper-triangular B_i's is easy, it does not immediately suggest an algorithm for matrix spaces in $\mathbf{UT}$.

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

Our main results are algorithms that run in polynomial time on an *algebraic RAM* [24], a random access machine in which the field operations as well as testing equality of field elements are performed at unit cost. Over finite fields, the straightforward implementations of these algorithms automatically have polynomial (in $\log \mathbb{F}$ and n) Boolean ("bit") complexity. With some effort, we are also able to present deterministic algorithms over the rationals which have Boolean complexity polynomial in the number of bits representing the input data. We now state our main theorems.

Theorem 1. *There are deterministic algorithms which solve the SMR on an algebraic RAM for $\mathbb{F}$ or over $\mathbb{Q}$, respectively, in polynomial time if $\mathcal{B}$ is spanned by rank-1 matrices. If the size of the base field is at least $n+1$, the algorithm solves the constructive SMR, and it also outputs a* $\mathrm{corank}(\mathcal{B})$*-singularity witness.*

Theorem 2. *Assume that the size of the base field $\mathbb{F}$ is at least $n+1$. Then there are deterministic polynomial-time algorithms which solve the constructive SDIT on an algebraic*

RAM for $\mathbb{F}$ or over $\mathbb{Q}$, respectively, if $\mathcal{B}$ is triangularizable. Furthermore, when $\mathcal{B}$ is singular, the algebraic RAM algorithm also outputs a singularity witness.

Theorem 1 can be slightly strengthened as follows: 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. See Remark 15 (2) for the work needed to achieve this.

Let us comment briefly on the framework for 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 [30]). These were originally defined in [33] 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 $\mathbf{UT}$ using the first Wong sequence. We remark that the application of the second Wong sequence is not new. Similar techniques were used in [22] to find maximum rank matrices in the case where rank one generators for $\mathcal{B}$ were given. Furthermore, while preparing the present version, we became aware of the paper [13] by Fortin and Reutenauer in which essentially the same method is used for testing existence of corank($\mathcal{B}$)-singularity witnesses (on a randomized algebraic RAM).

1.1 Comparison with previous works

The idea of singularity witnesses was already present in Lovász's work [26]. Among other things, Lovász showed that for the rank-1 spanned case, the equality $\operatorname{corank}(\mathcal{B}) = \operatorname{disc}(\mathcal{B})$ holds, by reducing it to Edmonds' Matroid Intersection theorem [10], which in turn can be deduced from Rado's matroidal generalization of Hall's theorem [28] (see also [32]). Inspired by this fact, Gurvits introduced the term *Edmonds-Rado property* for membership in the class of matrix spaces which are either nonsingular, or have a singularity witness. Throughout this paper we refer to this class as the *Edmonds-Rado class*. Gurvits listed several subclasses of the Edmonds-Rado class, including $\mathbf{R_1}$ (by the aforementioned result of Lovász) and $\mathbf{UT}$. A well-known example of a matrix space outside the Edmonds-Rado class is the linear space of skew symmetric matrices of size 3 [26].

As we stated already, Gurvits has presented a polynomial-time deterministic algorithm for the SDIT over subfields of $\mathbb{C}$ for matrix spaces in the Edmonds-Rado class. Therefore over these fields, his algorithm covers the SDIT for $\mathbf{R_1}$ and for $\mathbf{UT}$. Our algorithms (in the algebraic RAM model) are valid over arbitrary sufficiently large 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 27), so solving SMR for $\mathbf{UT}$ 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. 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 even over arbitrarily small finite fields, settling an open problem of Gurvits.

Over fields of constant size, the SMR has certain practical implications [19, 20], 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 ([19], building on [15, 27]), and when the mixed matrix is skew-symmetric and every variable appears at most twice ([16, 14]). Finally in [22], Ivanyos, Karpinski and Saxena present a deterministic polynomial-time algorithm for the case when among the input matrices $B_1, \ldots, B_m$ 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) [31, 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. [12] 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 [7].

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. Section 5 is devoted to the algorithm for triangularizable matrix spaces, proving Theorem 2. Finally, in Section 6 we propose and investigate some natural subclasses of the Edmonds-Rado class.

2 Wong sequences for pairs of matrix spaces

For $n \in \mathbb{N}$, we set $[n] = \{1, \ldots, n\}$. 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 framework for the algorithms in this work. Let V and V' be finite dimensional vector spaces over a field $\mathbb{F}$, and let $\mathrm{Lin}(V, V')$ be the vector space of linear maps from V to V'. Suppose $n = \dim(V)$ and $n' = \dim(V')$.

Let $U \leq V$ and $W \leq V'$ be subspaces of V and V', respectively. For $A \in \mathrm{Lin}(V, V')$, the image of U under A is $A(U) = \{A(u) \mid u \in U\}$, and the preimage of W under A is $A^{-1}(W) = \{v \in V \mid A(v) \in W\}$. To define generalized Wong sequences, the first step is to generalize the definitions of image and preimage under a single linear map A, to those under a matrix space $\mathcal{A} \leq \mathrm{Lin}(V, V')$.

Naturally, the image of U under $\mathcal{A}$ is the span of the images of U under every $A \in \mathcal{A}$, that is $\mathcal{A}(U) = \langle \cup_{A\in\mathcal{A}} A(U)\rangle = \langle \{A(u) \mid A \in \mathcal{A}, u \in U\}\rangle$. On the other hand, the preimage of W under $\mathcal{A}$ may be somewhat unexpected. It turns out that we need to take the intersection of the preimages of W under every $A \in \mathcal{A}$, that is $\mathcal{A}^{-1}(W) = \cap_{A\in\mathcal{A}} A^{-1}(W) = \{v \in V \mid \forall A \in \mathcal{A}, A(v) \in W\}$. Note that $\mathcal{A}(U)$ (resp. $\mathcal{A}^{-1}(W)$) is a subspace of V' (resp. V). Moreover, if $\mathcal{A}$ is spanned by $\{A_1, \ldots, A_m\}$, 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.

Lemma 3. *For $\mathcal{A}, \mathcal{B} \leq \mathrm{Lin}(V, V')$, 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 two Wong sequences for a pair of matrix subspaces.

Definition 4. *Let $\mathcal{A}, \mathcal{B} \leq \mathrm{Lin}(V, V')$. 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 $Ax - B$ [33, 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 Lemma 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 Lemma 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 $\ell \in \{0, \ldots, n\}$, such that $U_0 \supset U_1 \supset \cdots \supset U_\ell = U_{\ell+1} = \cdots$. 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 Lemma 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 $\langle , \rangle$. For a linear map $A : V \to V'$ let $A^{\mathrm{T}} : V' \to V$ stand for the transpose of A with respect to $\langle , \rangle$. This is the unique map with the property $\langle A^{\mathrm{T}}(u), v\rangle = \langle u, A(v)\rangle$, for all $u \in V'$ and $v \in V$. For a matrix space $\mathcal{A}$, let $\mathcal{A}^{\mathrm{T}}$ be the space $\{A^{\mathrm{T}} | A \in \mathcal{A}\}$. 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 $((\mathcal{A}^{\mathrm{T}})^{-1}(U))^\perp = \mathcal{A}(U^\perp)$, and $(\mathcal{A}^{\mathrm{T}}(V))^\perp = \mathcal{A}^{-1}(V^\perp)$. 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 $(\mathcal{A}^{\mathrm{T}}, \mathcal{B}^{\mathrm{T}})$, 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 $\mathcal{A}^{\mathrm{T}}(U^\perp)$ and finally $\mathcal{A}^{-1}(U) = (\mathcal{A}^{\mathrm{T}}(U^\perp))^\perp$. Therefore we have

Proposition 8. *Wong sequences can be computed in time using $(n+n')^{O(1)}$ on an algebraic RAM.*

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 $\mathrm{im}(A)$ 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 rank-1 spanned matrix spaces

3.1 Second Wong sequences and singularity witnesses

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 $A \in \mathrm{Lin}(V, V')$ we set $\mathrm{corank}(A) = \dim(\ker(A))$. For $\mathcal{B} \leq \mathrm{Lin}(V, V')$, the concepts of c-singularity witnesses, $\mathrm{disc}(\mathcal{B})$ and $\mathrm{corank}(\mathcal{B})$, defined for the case when $n = n'$, can be generalized naturally to $\mathcal{B}$. We also have that $\mathrm{corank}(\mathcal{B}) \geq \mathrm{disc}(\mathcal{B})$, and that a $\mathrm{corank}(\mathcal{B})$-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 $A \in \mathcal{B} \leq \mathrm{Lin}(V, V')$, and let W^* be the limit of the second Wong sequence of $(A, \mathcal{B})$. There exists a $\mathrm{corank}(A)$-singularity witness of $\mathcal{B}$ if and only if $W^* \subseteq \mathrm{im}(A)$. If this is the case, then A is of maximum rank and $A^{-1}(W^*)$ is a $\mathrm{corank}(\mathcal{B})$-singularity witness.*

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

Let us now suppose that some $U \leq V$ is a $\mathrm{corank}(A)$-singularity witness, that is $\dim(U) - \dim(\mathcal{B}(U)) \geq \mathrm{corank}(A)$. Then $\dim(U) - \dim(A(U)) \geq \mathrm{corank}(A)$ because $A \in \mathcal{B}$. Since the reverse inequality always holds without any condition on U, we have $\dim(U) - \dim(A(U)) = \mathrm{corank}(A)$. Similarly we have $\dim(U) - \dim(\mathcal{B}(U)) = \mathrm{corank}(A)$, 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 $\dim(S) - \dim(A(S)) = \mathrm{corank}(S)$ 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 $\mathrm{corank}(A)$-singularity witness obviously implies that A is of maximum rank, and when $W^* \subseteq \mathrm{im}(A)$ we have already seen that $A^{-1}(W^*)$ is a $\mathrm{corank}(A)$-singularity witness of $\mathcal{B}$. Since $\mathrm{corank}(A) = \mathrm{corank}(\mathcal{B})$, it is also a $\mathrm{corank}(\mathcal{B})$-singularity witness. □

We remark that in [13], a slightly different version of this statement is proved. We decided to keep our original proof for completeness. In our terminology, Theorem 3 of [13] states that the existence of a $\mathrm{corank}(A)$-singularity witness is equivalent to the equality $\dim(A^{-1}(W^*)) = \dim(W^*) + \dim(\ker(A))$. Both versions offer a straightforward method for testing existence of (and computing) $\mathrm{corank}(A)$-singularity witnesses. Besides that our version resembles the concept of augmenting paths in algorithms for matchings in bipartite graphs, it offers the possibility of stopping the construction of the Wong sequence at the point after which (while working over the rationals) data blowup can occur; this data blowup can occur if we adopt the naive way of computing the preimage of a subspace under A. Before that point, we will make use of a pseudo-inverse of A. 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 $A' : V' \to V$ is a *pseudo-inverse* of A if the restriction of A' to $\mathrm{im}(A)$ 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 $\mathrm{im}(A)$ in V'. Then take the map $A'_0 : \mathrm{im}(A) \to U$ such that AA'_0 is the identity of $\mathrm{im}(A)$ and take an arbitrary nonsingular linear map $A'_1 : U' \to \ker(A)$. Finally let A' be the direct sum of A'_0 and A'_1.

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

Proof. It follows from Lemma 9 that a $\mathrm{corank}(A)$-singularity witness exists if and only if $W_i \subseteq \mathrm{im}(A)$, for $i = 1, \ldots, n$. Observing that $(\mathcal{B}A')^i(\ker(AA')) \subseteq W_i$ for $i = 1, \ldots, n$, to prove the equivalence it is sufficient to show that if $(\mathcal{B}A')^i(\ker(AA')) \subseteq \mathrm{im}(A)$ for $i = 1, \ldots, n$ then $W_i = (\mathcal{B}A')^i(\ker(AA'))$ for $i = 1, \ldots, n$. 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 $W \leq \mathrm{im}(A)$ 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 $W_i \subseteq \mathrm{im}(A)$ and A' is the inverse of A on $\mathrm{im}(A)$.

Based on this equivalence, testing the existence of a $\mathrm{corank}(A)$-singularity witness can be accomplished by a simple algorithm. First compute a basis for $\mathcal{B}$, and then multiply it by A' to obtain a basis for $\mathcal{B}A'$. Compute also a basis for $\ker(AA')$. We now describe how to compute bases for the subspaces in the second Wong sequence until either we find i such that $W_i \not\subseteq \mathrm{im}(A)$ or we compute W^*. A basis for W_1 can be obtained by applying the basis elements of $\mathcal{B}A'$ to the basis elements of $\ker AA'$ and then selecting a maximal set of linearly independent vectors. Having computed a basis for W_i, we stop if it contains an element outside $\mathrm{im}(A)$. Otherwise we apply the basis elements of $\mathcal{B}A'$ to the basis elements of W_i, and select a maximal set of linearly independent vectors to obtain a basis for W_{i+1}. When $W_{i+1} = W_i$ we can stop since $W^* = W_i$.

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

3.2 The power overflow problem

For $A \in \mathcal{B} \leq \mathrm{Lin}(V, V')$, 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 $W^* \subseteq \mathrm{im}(A)$ 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 large enough fields, in 2-dimensional matrix spaces $\mathcal{B}$, the equality $\mathrm{corank}(\mathcal{B}) = \mathrm{disc}(\mathcal{B})$ holds.

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

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

Lemma 12. *Assume that $|\mathbb{F}| > n$. Let $A \in \mathcal{B} \leq \mathrm{Lin}(V, V')$, 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*

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

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 Fact 11 and Lemma 10, A is of maximum rank in $\langle A, B\rangle$ exactly when $(\langle B, A\rangle A')^i(\ker(AA')) \subseteq \mathrm{im}(A)$, for all $i \in [n]$. From that we can conclude since A' is the inverse of A on $\mathrm{im}(A)$. □

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 on an algebraic RAM as well as over $\mathbb{Q}$.*

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 $(\mathcal{B}A')^i(\ker(AA')) \subseteq \mathrm{im}(A)$ 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 $(BA')^i(\ker(AA')) \not\subseteq \mathrm{im}(A)$. We apply the algorithm of Theorem 14 with input $\mathcal{B}A'$, $\ker(AA')$ and $\mathrm{im}(A)$, 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 of $A+\lambda B$ 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 $\lambda_1, \ldots, \lambda_{n+1}$, and we know that for some $1 \leq j \leq n+1$ we have $\mathrm{rank}(A+\lambda_j B) > \mathrm{rank}(A)$. We replace A by $A+\lambda_j B$ and restart the process.

Over $\mathbb{Q}$, at the end of each iteration, by a reduction procedure described in [8] we can achieve that the matrix A, written as a linear combination of $B_1, \ldots, B_m$ has coefficients from a fixed subset $K \subseteq \mathbb{Q}$ of size $n+1$ (say, $K = \{0, \ldots, n\}$). In fact, if $A = \alpha_1 B_1 + \alpha_2 B_2 \ldots + \alpha_m B_m$ has rank r then for at least one $\kappa_1 \in K$ the matrix $\kappa_1 B_1 + \alpha_2 B_2 \ldots + \alpha_m B_m$ 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 maximal 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. This finishes the case for $|\mathbb{F}| \geq n+1$.

When $|\mathbb{F}| < n+1$, we can compute the maximum rank 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 $\mathrm{corank}(\mathcal{B}) = \mathrm{disc}(\mathcal{B})$ holds over any field if $\mathcal{B}$ is spanned by an arbitrary matrix and by rank one matrices, see [22]. □

Remark 15. *As mentioned in the introduction, we can generalize to the setting when $\mathcal{B}$ is spanned by rank-1 matrices and an arbitrary matrix, as follows. Let $\mathcal{B}'$ be the subspace of $\mathcal{B}$ generated by rank-1 matrices. As indicated in Lemma 18 in the next section, in this case the algorithm for power overflow problem with special $U = \ker(AA')$ and $U' = \mathrm{im}(A)$ is still guaranteed to succeed if $A \notin \mathcal{B}'$. Furthermore, in the update step, the resulting matrix of higher rank keeps the property of not in $\mathcal{B}'$. So from the given basis $B_1, \ldots, B_m$ for $\mathcal{B}$, we apply the procedure in Theorem 1 using B_i as the starting point, for each B_i. Then it is ensured that for those $B_i \notin \mathcal{B}'$ this procedure will succeed in finding a matrix with maximal rank. Otherwise if $B_i \in \mathcal{B}'$, then power overflow problem will either get* **no**, *or detect that $\mathcal{H}_\ell \cdots \mathcal{H}_1(U) \subseteq U'$, so ending with a safe return.*

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

In this section we prove **Theorem 14**.

The setting. Given subspaces U, U' of $\mathbb{F}^n$ as well as a basis $\{D_1, \ldots, D_m\}$ 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 $\mathcal{T}_1 = \{D_1, \ldots, D_m\}$. 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 form the products of elements in $\mathcal{T}_j$ and $\mathcal{T}_1$.

Helpful subspaces of $\mathcal{D}$. Recall that our goal is to find D such that $D^\ell(U) \not\subseteq U'$. To achieve this goal, for $i \in [\ell]$, we define subspaces $\mathcal{H}_i$ of $\mathcal{D}$, which play a crucial role in the algorithm:

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

Let us examine the meaning for some matrix X to be in $\mathcal{H}_i$. Let P be a product of ℓ elements from $\mathcal{D}$, and suppose X appears in P. Then $X \in \mathcal{H}_i$ implies that, as long as X appears in P at the jth position, $j \neq i$, then it must be that $P(U) \subseteq U'$. In other words, for P to be able to pull U out of U', it is necessary that X appears at the ith position.

The following lemma explains why $\mathcal{H}_i$'s are useful for the purpose of powerflow problem.

Lemma 16. *For a matrix $X = X_1 + \ldots + X_\ell$ 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 $X^\ell = \sum_\sigma X_{\sigma(\ell)} \cdots X_{\sigma(1)}$, where the summation is over the maps $\sigma : [\ell] \to [\ell]$. 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 $\mathcal{H}_{\sigma(j)}$. □

Furthermore, $\mathcal{H}_i$'s can be computed efficiently on an algebraic RAM as well as over $\mathbb{Q}$ as follows. Let $x_1, \ldots, x_m$ 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 $\{D_1, \ldots, D_m\}$ to basis elements for U.

Back to rank-1 spanned setting. In general, $\mathcal{H}_i$ can be 0. In our setting, due to the existence of a basis of rank-1 matrices, fortunately this is far from the case.

Lemma 17. *Suppose $\mathcal{B}$ is rank-1 spanned, and ℓ is the smallest integer such that $\mathcal{D}^\ell(U) \not\subseteq U'$. Then the following hold: (1) $\forall i \in [\ell]$, $\mathcal{H}_i \neq 0$; (2) $\mathcal{H}_\ell \cdots \mathcal{H}_1(U) \not\subseteq U'$.*

Proof. Assume that $\mathcal{D}$ is spanned by the rank one matrices $C_1, \ldots, C_m$, where C_i may be over an extension field $\mathbb{F}'$ of $\mathbb{F}$.

Let us first consider the case when C_i's are matrices over $\mathbb{F}$. Then there exist indices $k_1, \ldots, k_\ell$ such $C_{k_\ell} \cdots C_{k_1}(U) \not\subseteq U'$. We show that $C_{k_i} \in \mathcal{H}_i$, for $i \in [\ell]$, proving (1). This also implies immediately $\mathcal{H}_\ell \cdots \mathcal{H}_1(U) \not\subseteq U'$, proving (2).

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 generalize to C_i's over an extension field $\mathbb{F}'$, it suffices to lift all objects ($\mathcal{D}$, $\mathcal{H}_i$, U and U') to their spans with the extension field $\mathbb{F}'$ (denoted as $\mathbb{F}'\mathcal{D}$, $\mathbb{F}'\mathcal{H}_i$, $\mathbb{F}U$ and $\mathbb{F}U'$). After going through the above argument, we have $\mathbb{F}'\mathcal{H}_i \neq 0$ and $\mathbb{F}'\mathcal{H}_\ell \cdots \mathbb{F}'\mathcal{H}_1(\mathbb{F}'U) \not\subseteq \mathbb{F}'U'$. We then have $\mathcal{H}_i \neq 0$ and $\mathcal{H}_\ell \cdots \mathcal{H}_1(U) \not\subseteq U'$, as it is not hard to see that $\mathbb{F}'\mathcal{H}_i$, and $\mathbb{F}'\mathcal{H}_\ell \cdots \mathbb{F}'\mathcal{H}_1(\mathbb{F}'U)$, are spans of $\mathcal{H}_i$ and $\mathcal{H}_\ell \cdots \mathcal{H}_1(U)$ with the extension field $\mathbb{F}'$. □

That is, in out setting, not only $\mathcal{H}_i \neq 0$, but $\mathcal{H}_\ell \cdots \mathcal{H}_1$ is able to pull U outside U' (instead of using the full power of $\mathcal{D}^\ell$).

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 $\mathcal{H}_\ell$ 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 17 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 $D = X_1 + \ldots + X_\ell$, then by Lemma 16 we have $D^\ell(U) \not\subseteq U'$. We return D and ℓ. This finishes the proof of Theorem 14.

Finally, we introduce the following slight extension of Lemma 17 for special subspaces U, U', as applicable to Remark 15 (2).

Lemma 18. *Assume that $\mathcal{D}$ is spanned by rank one matrices and a projection to U' having kernel U. Then $\mathcal{H}_\ell \cdots \mathcal{H}_1 U \not\subseteq U'$.*

Proof. Identical with the proof of Lemma 17, based on the observation that a projection with the prescribed properties can be deleted from any product mapping U outside U'. □

5 The first Wong sequence and triangularizable matrix spaces

5.1 The connection

To tackle the triangularizable matrix spaces, our starting point is the following lemma, which connects first Wong sequences with singularity witnesses.

Lemma 19. *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.

Proof. If $\dim(U^*) > \dim(\mathcal{B}(U^*))$ then U^* is a singularity witness. If $\dim(U^*) = \dim(\mathcal{B}(U^*))$ then the choice of P and Q corresponds to an appropriate basis change transformation. To see that $\mathcal{B}$ is nonsingular in the X-block, note that $A \in \mathcal{B}$ and $A(U^*) = \mathcal{B}(U^*)$. □

Lemma 19 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 20. *Let $\mathcal{B} \leq \mathbb{F}$ be given by a basis $\{B_1, \ldots, B_m\}$, 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 $\cap_{i\in[m]} \ker(B_i) \neq 0$, or $\exists j \in [m]$ 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 $\dim(U) = \dim(B_j(U))$. Let $B_i^* : \mathbb{F}^n/U \to \mathbb{F}^n/\mathcal{B}(U)$ be the linear map induced by B_i, for $i \in [m]$. Then $\mathcal{B}^* = \langle B_1^*, \ldots, B_m^* \rangle$ 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 $0 = V_0 \subset V_1 \subset \cdots \subset V_n = V$. For $\mathcal{A} \leq \mathrm{Lin}(V, V')$ with $\dim(V) = \dim(V') = n$, the matrix space $\mathcal{A}$ is triangularizable if and only if $\exists$ complete flags $0 = V_0 \subset V_1 \subset \cdots \subset V_n = V$ and $0 = V_0' \subset V_1' \subset \cdots \subset V_n' = V'$ 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 $S_i = \langle c_1, \ldots, c_i \rangle$ and $T_i = \langle d_1, \ldots, d_i \rangle$. 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 $S_0^* \subseteq \cdots \subseteq S_n^*$ and $T_0^* \subseteq \cdots \subseteq T_n^*$. 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 $\dim(\mathbb{F}'U) = \dim(B_j(\mathbb{F}'U))$, $\dim(B_j(S_i \cap \mathbb{F}'U)) \geq \dim(S_i \cap \mathbb{F}'U)$. 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 $S_0^* \subset S_{j_1}^* \subset \cdots \subset S_{j_\ell}^* = S_n^*$ and $T_0^* \subset T_{k_1}^* \subset \cdots \subset T_{k_\ell}^* = T_n^*$, 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 $\mathbb{F}'$. □

5.2 An algorithm on an algebraic RAM

Suppose we are given a basis $\{B_1, \ldots, B_m\}$ for $\mathcal{B} \leq M(n, \mathbb{F})$ which is triangularizable over an extension field $\mathbb{F}'$ of $\mathbb{F}$, i. e., $B_i = DB_i'C^{-1}$ for some nonsingular $C, D \in M(n, \mathbb{F}')$, and $B_i' \in M(n, \mathbb{F}')$ is upper triangular for every $i \in [m]$. Our problem is to determine whether there exists a nonsingular matrix in $\mathcal{B}$ or not and finding such a matrix if exists.

Given the preparation of Lemma 20, here is the outline of an algorithm using polynomially many arithmetic operations. The algorithm recurses on the size of the matrices, with the base case being the size 1. It checks at the beginning whether $\cap_{i\in[m]} \ker(B_i) = 0$. If this is the case then it returns $\cap_{i\in[m]} \ker(B_i)$ 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 20 (1) there exists $j \in [m]$ such that $U_j^* \neq 0$ and $B_j(U) = \mathcal{B}(U)$. The algorithm then recurses on the induced actions B_i^*'s of B_i's, which are also triangularizable by Lemma 20 (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.

We expand the above idea into a rigorous algorithm, called TriAlgo and present it in Algorithm 1. This algorithm requires polynomially many arithmetic operations, and therefore of polynomial complexity in finite fields. The input of the algorithm can be an arbitrary matrix space (not necessarily triangularizable), but it may fail in certain cases. For triangularizable matrix spaces the algorithm would not fail due to Lemma 20. Note that though the algorithm works assuming triangularizability over some extension field, the algorithm itself does not need to deal with the field extension explicitly by Lemma 20, given that $\mathbb{F}$ is large enough. To allow for recursion, the output of the algorithm can be one of the following: the first is an explicit linear combination of the given matrices, which gives a nonsingular matrix. The second is a singular subspace witness. The third one is `Fail`.

Regarding implementation, it might be needed to comment on Line 13. At this point we have that $\dim(U^*) \leq \dim(\mathcal{B}(U^*)) \leq \dim(B_j(U^*)) \leq \dim(U^*)$. Thus $\dim(\mathcal{B}(U^*)) = \dim(U^*)$, and note that $B_i(U^*) \subseteq \mathcal{B}(U^*)$, for all $i \in [m]$. Then two bases of $\mathbb{F}^n$ can be formed by extending bases of U^* and $\mathcal{B}(U^*)$ respectively, and w.r.t. these two bases the induced action B_i from $\mathbb{F}^n/U^*$ to $\mathbb{F}^n/\mathcal{B}(U^*)$ can be read off easily.

For correctness we distinguish among the types of output of the algorithm, and show that they indeed have the required property.

If $(\alpha_1, \ldots, \alpha_m)$ is returned: This case occurs in Line 2 and Line 20. Line 2 is trivial. If the algorithm reaches Line 20, we claim that there exists $(\lambda, \mu) \in \Lambda \times \Lambda$ s.t. $\lambda B_j + \mu E$ is nonsingular. Let P and Q be the matrices from Lemma 19. Thus $\forall i \in [m]$, QB_iP^{-1} is of the form: $\begin{bmatrix} X_i & Y_i \\ 0 & Z_i \end{bmatrix}$, where X_i is of size $(n-\ell)\times(n-\ell)$ and Z_i is of size ℓ by ℓ. As X_j is nonsingular and $\sum_{i\in[m]} \alpha_i Z_i$ is nonsingular, $\det(xB_j + yE)$ is a nonzero polynomial, thus from Schwartz-Zippel lemma the existence of (λ, μ) in $\Lambda \times \Lambda$ is ensured.

If a subspace of $\mathbb{F}^n$ is returned: This case occurs in Line 2, 4, 8 and 16. All are straightfor-

Algorithm 1: TriAlgo$(B_1, \ldots, B_m)$

Input: $\mathcal{B} = \langle B_1, \ldots, B_m \rangle \subseteq M(n, \mathbb{F})$.

Output: One of the following: (1) $(\alpha_1, \ldots, \alpha_m) \in \mathbb{F}^m$ s.t. $\sum_{i\in[m]} \alpha_i B_i$ is nonsingular. (2) A singular subspace witness $U \leq \mathbb{F}^n$. (3) `Fail`.

`// Base case`

if $n = 1$ **then**

  If $\exists$ nonzero B_i, **return** $(0, \ldots, 1, \ldots, 0)$ where 1 is at the ith position. Otherwise **return** $\mathbb{F}^n$.

`// Start of the recursive step.`

`// If` $\cap_{i\in[m]} \ker(B_i) \neq 0$ `then` $\cap_{i\in[m]} \ker(B_i)$ `is a singular witness itself.`

if $\cap_{i\in[m]} \ker(B_i) \neq 0$ **then**

  return $\cap_{i\in[m]} \ker(B_i)$.

forall the $i \in [m]$ **do**

  $U_i^* \leftarrow$ the limit of the first Wong sequence of $(B_i, \mathcal{B})$.

if $\exists i \in [m]$, $\dim(\mathcal{B}(U_i^*)) < \dim(U_i^*)$ **then**

  return U_i^*

if $\nexists j$ *s.t.* $\dim(U_j^*) > 0$ **then**

  return `Fail`

$U^* \leftarrow U_j^*$ where U_j^* satisfies that $\dim(U_j^*) > 0$;

forall the $i \in [m]$ **do**

  $B_i^* \leftarrow$ the induced linear map of B_i from $\mathbb{F}^n/U^*$ to $\mathbb{F}^n/\mathcal{B}(U^*)$.

`// Recursive call.`

$X \leftarrow$ TriAlgo$(B_1^*, \ldots, B_m^*)$;

if X *is a singular subspace witness* W/U^* **then**

  return the full preimage of W/U^* in the canonical projection $\mathbb{F}^n \to \mathbb{F}^n/U^*$.

else if X *is* $(\alpha_1, \ldots, \alpha_m)$ **then**

  $\Lambda \leftarrow$ a set of field element of size $n + 1$;

  $E \leftarrow \sum_{i\in[m]} \alpha_i B_i$;

  Choose $(\lambda, \mu) \in \Lambda \times \Lambda$, s.t. $\lambda B_j + \mu E$ is nonsingular;

  return $(\mu\alpha_1, \ldots, \mu\alpha_{j-1}, \mu\alpha_j + \lambda, \mu\alpha_{j+1}, \ldots, \mu\alpha_m)$

else if $X =$ `Fail` **then**

  return `Fail`

ward.

The case of `Fail`: After Line 3 $\cap_{i\in[m]} \ker(B_i) = 0$. Then Lemma 20 ensures that `Fail` cannot be returned for triangularizable matrix spaces.

5.3 An algorithm over the rationals

To obtain a polynomial-time algorithm over rationals, we give first a characterization of triangularizability of a nonsingular matrix space.

Lemma 21. *Assume $\mathcal{B} \leq M(n,\mathbb{F})$ contains a nonsingular matrix S. Then $\mathcal{B}$ is triangularizable over $\mathbb{F}$ if and only if there exists a nonsingular matrix $D \in M(n,\mathbb{F})$ such that $D^{-1}\mathcal{B}S^{-1}D$ consists of upper triangular matrices.*

Proof. $\Rightarrow$: Assume that $D^{-1}\mathcal{B}C$ consists of upper triangular matrices. Then $C^{-1}S^{-1}D = (D^{-1}SC)^{-1}$ is upper triangular as well, whence – as products of upper triangular matrices remain upper triangular – $D^{-1}\mathcal{B}S^{-1}D = (D^{-1}\mathcal{B}C)(C^{-1}S^{-1}D)$ also consists of upper triangular matrices.
$\Leftarrow$: Assume that $D^{-1}\mathcal{B}S^{-1}D$ consists of upper triangular matrices. Put $C = S^{-1}D$. □

We have the following criterion of triangularizability:

Lemma 22. *Let $\mathcal{A} \leq M(n,\mathbb{F})$ containing the identity matrix and let $\mathbb{F}'$ be the algebraic closure of $\mathbb{F}$. Then there exists $D \in M(n,\mathbb{F}')$ such that $D^{-1}\mathcal{A}D$ consists of upper triangular matrices (over $\mathbb{F}'$) if and only if*

$$(\mathcal{A}^{n^2}[\mathcal{A},\mathcal{A}]\mathcal{A}^{n^2})^n = (0).$$

Here $[\mathcal{A},\mathcal{A}]$ is the space spanned by the commutators $[X,Y] = XY - YX$ $(X,Y \in \mathcal{A})$.

Proof. Put $\mathcal{D} = \mathcal{A}^{n^2}$. Then $\mathcal{D}$ is the matrix algebra generated by $\mathcal{A}$. The formula expresses that the two-sided ideal of $\mathcal{D}$ generated by the commutators from $\mathcal{A}$ is nilpotent. Let $\mathcal{D}' = \mathbb{F}' \otimes \mathcal{D}$. Then the formula is also equivalent to that the ideal of $\mathcal{D}'$ generated by the commutators is nilpotent. This is further equivalent to that the factor algebra $\mathcal{D}'/Rad(\mathcal{D}')$ is commutative. However, over an algebraically closed field a matrix algebra is nilpotent if it is a conjugate of a subalgebra of the upper triangular matrices. (To see one direction, observe that the whole algebra of the upper triangular matrices and hence every subalgebra of it has this property. As for reverse implication, note that all the irreducible representations of an algebra over an algebraically closed field which is commutative by its radical are one-dimensional and hence a composition series gives a complete flag consisting of invariant subspaces.) □

Corollary 23. *Assume that we are given a nonsingular $S \in \mathcal{B}$. Then there is a polynomial time algorithm (on an algebraic RAM as well as the case $\mathbb{F} = \mathbb{Q}$) which decides whether or not there exists an extension of $\mathbb{F}$ over which $\mathcal{B}$ is triangularizable.*

Again, we actually have an algorithm using a polynomial number of arithmetic operations and equality tests in the black box model for $\mathbb{F}$.

With these preparations we are now ready to prove Theorem 2.

Proof of Theorem 2. On an algebraic RAM Algorithm 1 is all we need. Over rationals we shall perform a reduction to finite fields via Lemma 22.

We assume that $\mathcal{B}$ is given by matrices $B_1, \ldots, B_m$ over $\mathbb{Q}$. Multiplying by a common denominator for the entries, we can achieve the situation when the entries of $B_1, \ldots, B_m$ are integers. Let b be a bound on the of absolute values of the entries of $B_1, \ldots, B_m$. Then a polynomial-time algorithm should run in time polynomial in n and $\log b$. If $\mathcal{B}$ is nonsingular then there exist integers $\lambda_1, \ldots, \lambda_m$, each between 0 and n such that $S = \lambda_1 B_1 + \ldots + \lambda_m B_m$ is nonsingular. The absolute value of the determinant of S is a nonzero integer whose logarithm is bounded by a polynomial in $\log b$ and n. It follows that there is a prime p bounded by an (explicit) polynomial in $\log b$ and n that does not divide the determinant of S.

Let $S' = \det(S)S^{-1}$. We reduce the problem modulo p. We see that S and S' are integral matrices and both are invertible module p. Furthermore, if $\mathcal{B}$ is triangularizable over an extension of $\mathbb{Q}$, by Lemma 22 all length-n products of elements of the form $B_{i_1}S' \cdots B_{i_{n^2}}S'[B_{j_1}S', B_{j_2}S']B_{k_1}S' \cdots B_{k_{n^2}}S'$ vanish, and this will be the case modulo p as well. It follows that the subspace of matrices over $\mathbb{F}_p$, spanned by the matrices B_iS', reduced modulo p, can be triangularized over an extension field of $\mathbb{F}_p$. But then the space spanned by B_i is also triangularizable (over the same extension).

Thus if $\mathcal{B}$ is nonsingular and triangularizable over an extension of $\mathbb{Q}$ then there is a prime p greater than n but smaller than the value of an explicit polynomial function in $\log b$ and n, such that the reduction modulo p gives a nonsingular matrix space which is triangularizable over an extension field of $\mathbb{F}_p$. The algorithm consists of taking the primes p up to the polynomial limit and applying the generic method over $\mathbb{F}_p$ to the reduced setting. The method either finds a p and an integer combination of $B_1, \ldots, B_m$ which is nonsingular even modulo p or, concludes that $\mathcal{B}$ cannot be nonsingular and triangularizable at the same time. □

6 On the Edmonds-Rado class and some subclasses

6.1 Matrix spaces not in the Edmonds-Rado class

Recall that $\mathcal{B} \leq M(n, \mathbb{F})$ is in the Edmonds-Rado class if either $\mathcal{B}$ contains nonsingular matrices, or $\mathcal{B}$ is singular and there exists a singularity witness of $\mathcal{B}$. Recall that in Section 1 we defined the discrepancy $\mathrm{disc}(\mathcal{B}) = \max\{c \in \mathbb{N} \mid \exists\ c\text{-singularity witness of } \mathcal{B}\}$, and from the definition it is clear that $\mathrm{corank}(\mathcal{B}) \geq \mathrm{disc}(\mathcal{B})$. In terms of discrepancy, $\mathcal{B}$ is in the Edmonds-Rado class, if $\mathrm{disc}(\mathcal{B}) = 0 \iff \mathrm{corank}(\mathcal{B}) = 0$.

A well-known example of a matrix space not in the Edmonds-Rado class is the class sk_3

of 3 dimensional skew symmetric matrices, generated for example by the following:

$$\mathrm{sk}_3 = \left\langle \begin{bmatrix} 0 & 1 & 0 \\ -1 & 0 & 0 \\ 0 & 0 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 0 \\ 0 & 0 & 1 \\ 0 & -1 & 0 \end{bmatrix}, \begin{bmatrix} 0 & 0 & 1 \\ 0 & 0 & 0 \\ -1 & 0 & 0 \end{bmatrix} \right\rangle.$$

Consider the following problem: if a matrix space $\mathcal{B}$ has a basis consisting of matrices with certain property, does it imply that $\mathcal{B}$ is in the Edmonds-Rado class? Gurvits has observed that if $\mathcal{B}$ has a basis consisting of triangular or semidefinite matrices then it is in the Edmonds-Rado class. We now show that the other two basis properties, namely consisting of projections or positive matrices, do not necessarily imply that $\mathcal{B}$ is in the Edmonds-Rado class. Recall that a matrix over $\mathbb{R}$ is positive if every entry in it is positive.

Let $\mathcal{B} = \langle B_1, \ldots, B_m \rangle \leq M(n, \mathbb{F})$, and let $A \in M(n, \mathbb{F})$ be an arbitrary nonsingular matrix. For $i \in [m]$, we define $Y_i = \begin{bmatrix} A & B_i \\ 0 & 0 \end{bmatrix}$ and $Z = \begin{bmatrix} 0 & 0 \\ A & 0 \end{bmatrix}$, and let $\mathcal{A} = \langle Y_1, \ldots, Y_m, Z \rangle$.

Lemma 24. *We have* $\mathrm{disc}(\mathcal{A}) = \mathrm{disc}(\mathcal{B})$.

Proof. Let $E_1 \leq \mathbb{F}^{2n}$ be the coordinate subspace generated by the first n coordinates, and $E_2 \leq \mathbb{F}^{2n}$ be the coordinate subspace generated by the last n coordinates.

To show that $\mathrm{disc}(\mathcal{A}) \geq \mathrm{disc}(\mathcal{B})$, we take a $\mathrm{disc}(\mathcal{B})$-discrepancy witness U of $\mathcal{B}$, and embed U into E_2. Then $\langle E_1 \cup U \rangle$ is a $\mathrm{disc}(\mathcal{B})$-singularity witness of $\mathcal{A}$.

To show that $\mathrm{disc}(\mathcal{A}) \leq \mathrm{disc}(\mathcal{B})$, let W be $\mathrm{disc}(\mathcal{A})$-singularity witness of $\mathcal{A}$. Let $W_1' = W' \cap E_1$ and $W_2' = W' \cap E_2$. Due to the form of the Y_i's and Z, $W' = W_1' \oplus W_2'$. In particular note that $W_2' = Z(W)$, and A is nonsingular. So if we set $R = \{w \in W \mid Z(w) = 0\}$, we have $R \leq E_2$, $\dim(R) = \dim(W) - \dim(W_2')$, and $\dim(W_1') \geq \dim(\mathcal{B}(R))$. Thus $\mathrm{disc}(\mathcal{A}) = \dim(W) - \dim(W') = (\dim(W) - \dim(W_2')) - \dim(W_1') \leq \dim(R) - \dim(\mathcal{B}(R)) \leq \mathrm{disc}(\mathcal{B})$. □

Proposition 25. *There exist matrix spaces generated by projections or positive matrices outside the Edmonds-Rado class*

Proof. For $i \in [m]$, we define $Y_i', Z' \in M(2n, \mathbb{F})$ by $Y_i' = \begin{bmatrix} A & B_i + A \\ 0 & 0 \end{bmatrix}$ and $Z' = \begin{bmatrix} 0 & 0 \\ A & A \end{bmatrix}$, and let $\mathcal{A}' = \langle Y_1', \ldots, Y_m', Z' \rangle$. It is easy to see that the Y_i''s and Z' can be obtained from the Y_i's and Z via simultaneous row and column operations. Note that simultaneous row and column operations do not change the rank or the discrepancy of a space, that is $\mathrm{corank}(\mathcal{A}') = \mathrm{corank}(\mathcal{A})$ and $\mathrm{disc}(\mathcal{A}') = \mathrm{disc}(\mathcal{A})$. Observe that $\mathrm{corank}(\mathcal{A}) = \mathrm{corank}(\mathcal{B})$, and by Lemma 24 we have $\mathrm{disc}(\mathcal{A}) = \mathrm{disc}(\mathcal{B})$. Therefore taking some $\mathcal{B}$ not in the Edmonds-Rado class (for example sk_3) it follows that $\mathcal{A}$ and $\mathcal{A}'$ are not in the Edmonds-Rado class. To finish the proof just note that if $A = I$, then Y_i''s and Z' are projections, and if A is a positive matrix with entries at least the absolute values of the entries in the B_i's, then Y_i''s and Z' are positive matrices. □

6.2 Compression spaces

If maxrk($\mathcal{B}$) is of primary interest, in analogy with the Edmonds-Rado class we can define the following matrix class. Here we allow non-square matrices from $M(n \times n', \mathbb{F})$. Recall that for $A \in M(n \times n', \mathbb{F})$, its rank is $\dim(\mathrm{im}(A))$, its corank is $\dim(\ker(A))$, and A is nonsingular if $\mathrm{rank}(A) = \min(n, n')$.

Following the terminology used in [11, 13], we call a matrix space $\mathcal{B} \leq M(n \times n', \mathbb{F})$ a *compression space* if $\mathcal{B}$ possesses corank($\mathcal{B}$)-singularity witnesses. In terms of discrepancy, $\mathcal{B}$ is a compression space if $\mathrm{corank}(\mathcal{B}) = \mathrm{disc}(\mathcal{B})$. Thus, by the result of Lovász discussed in Subsection 1.1, and by the result of Atkinson and Stephens used in Subsection 3.2, rank-one spanned matrix spaces as well as two-dimensional matrix spaces over sufficiently large base fields are compression spaces. As to Wong sequences, from Lemma 9 we immediately have that if $\mathcal{B} \leq M(n \times n', \mathbb{F})$ is a compression space, then for any $A \in \mathcal{B}$, A is of maximum rank if and only if the limit of the second Wong sequence of $(A, \mathcal{B})$ is contained in $\mathrm{im}(A)$.

It is clear that when $n = n'$, if $\mathcal{B}$ is a compression space then it is in the Edmonds-Rado class. The converse is not true.

Proposition 26. *There exists a matrix space in the Edmonds-Rado class which is not a compression space.*

The proof of Proposition 26 relies on the following lemma, which also explains why we do not expect to achieve rank maximization for upper triangular matrices in Theorem 2.

Lemma 27. *Rank maximization of matrix spaces can be reduced to rank maximization of matrix spaces with a basis of pairwise commuting, and strictly upper triangular matrices.*

Proof. For $\mathcal{B} = \langle B_1, \ldots, B_m \rangle \leq M(n \times n', \mathbb{F})$ we first pad 0's to make it a matrix space of $M(\max(n, n'), \mathbb{F})$. Then consider the matrix space in $M(2 \cdot \max(n, n'), \mathbb{F})$ generated by $C_1, \ldots, C_m$ where $C_i = \begin{bmatrix} 0 & B_i \\ 0 & 0 \end{bmatrix}$. □

Proof of Proposition 26. Consider the following matrix space: apply the construction in Lemma 27 with sk_3, and let the resulting matrix space be $\mathcal{B} \leq M(6, \mathbb{Q})$. $\mathcal{B}$ is in the Edmonds-Rado class as it is spanned by upper-triangular matrices. On the other hand $\mathcal{B}$ is not in the Edmonds-Rado class as $\mathrm{corank}(\mathcal{B}) = 4$ while $\mathrm{disc}(\mathcal{B}) = 3$. □

6.3 The black-box Edmonds-Rado class

Definition 28. *Let $\mathcal{B} \leq M(n \times n', \mathbb{F})$. $\mathcal{B}$ is in the black-box Edmonds-Rado class if the following two conditions hold: (1) there exists a* corank($\mathcal{B}$)*-singularity witness; (2) for any $A \in \mathcal{B}$, either A is of maximum rank, or $\mathcal{B}(\ker(A)) \not\subseteq \mathrm{im}(A)$.*

By the first condition, the black-box Edmonds-Rado class is a subclass of the compression spaces. Also note that $\mathcal{B}(\ker(A))$ is just the first item in the second Wong sequence of $(A, \mathcal{B})$. The second condition says that if A is non-maximum rank then already the first item in the

second Wong sequence excludes existence of corank(A)-singularity witnesses. In this case for any matrix B from $\mathcal{B}$ with $B(\ker(A)) \not\subseteq \mathrm{im}(A)$, we have $\mathrm{rank}(B) > \mathrm{rank}(A)$. Therefore in matrix spaces in this class the following simple algorithm finds an element of maximum rank over sufficiently large base fields.

Proposition 29. *Let $\mathcal{B} \leq M(n, \mathbb{F})$ be in the black-box Edmonds-Rado class, and assume $|\mathbb{F}| = \Omega(n)$. Then there exists a deterministic algorithm that solves the constructive SMR for $\mathcal{B}$ using polynomial number of arithmetic operations.*

Proof. Given $A \in \mathcal{B}$, we compute the rank of $A + \lambda B$ where λ is from a subset of $\mathbb{F}$ of size $\mathrm{rank}(A) + 1$ and B is from a basis of $\mathcal{B}$. If none of these matrices have rank larger than A, conclude that A is of maximum rank. Otherwise replace A with an $A + \lambda B$ of larger rank. Iterate the above procedure to obtain $A \in \mathcal{B}$ of maximum rank. □

As a justification for the name of the subclass, observe that this algorithm does not make use of any properties of matrices other that their rank. It even works in the setting that instead of inputting the basis $B_1, \ldots, B_m$ explicitly, we only know m and have an oracle which, on input $(\alpha_1, \ldots, \alpha_m)$ returns the rank of $\alpha_1 B_1 + \ldots + \alpha_m B_m$.

6.3.1 Some matrix spaces in the black-box Edmonds-Rado class.

While this class seems quite restrictive, it contains some interesting cases.

A first example is when $\mathcal{B}$ has a basis of positive semidefinite matrices. Let $\mathcal{B} = \langle B_1, \ldots, B_m \rangle \leq M(n, \mathbb{R})$, where B_i's are positive semidefinite. Then it is seen easily that A is of maximum rank if and only if $\ker(A) = \cap_{i \in [m]} \ker(B_i)$. In particular if A is not of maximum rank then there exists $v \in \ker(A)$ such that $B_j(v) \notin \mathrm{im}(A)$, for some $j \in [m]$.

Another more interesting scenario is from [6] (see also [22], Lemma 4.2). Let G be a finite dimensional associative algebra over $\mathbb{F}$ and let V, V' be semisimple G-modules. Let $\mathcal{B} = \mathrm{Hom}_G(V, V')$. Recall that a semisimple module is the direct sum of simple modules and that in a semisimple module every submodule has a direct complement. We know that $A \in \mathcal{B}$ is of maximum rank if and only if for every isomorphism type S of simple modules for A, the multiplicity of S in $\mathrm{im}(A)$ is the minimum of the multiplicities of S in U and V.

If A is not of maximum rank, then for some simple module S there is an isomorphic copy S_1 of S in $\ker(A)$ and there is a copy S_2 of S in V' intersecting $\mathrm{im}(A)$ trivially. Also, there are nontrivial homomorphisms mapping the first copy of S to the second one. For instance, any isomorphism $S_1 \to S_2$ can be extended to a homomorphism $V \to V'$ by the zero map on a direct complement of S_1.

On the other hand, if A is of maximum rank then for every simple submodule in $\ker(A)$, the copies in V' isomorphic to it are in $\mathrm{im}(A)$, therefore no simple constituent can be moved out of $\mathrm{im}(A)$ via the second Wong sequence.

7 Concluding remarks

Our main results are deterministic polynomial time algorithms for the *constructive* version of Edmond's problem (that is, finding nonsingular matrices) in certain subclasses of the Edmonds-Rado class. In the light of Gurvits' result on the non-constructive version, probably the most interesting open problem is the deterministic complexity of the constructive version for the whole Edmonds-Rado class. Regarding the Boolean complexity of some of our algorithms, the bottleneck is our limited knowledge about the possible blowup of the sizes of bases for the Wong sequences. We are not even aware of any good bound on the size of bases for singularity witnesses (except for the rank one generated case). In particular, we do not know the Boolean complexity of finding singularity witnesses for singular triangularizable matrix spaces over the rationals.

The deterministic or randomized complexity of finding rank one matrices spanning a rank-one generated space is another question which is open to our knowledge. We believe that the problem is hard. In contrast, triangularizing a triangularizable matrix space may be easier. In the special case when the space is triangularizable over the base field $\mathbb{F}$ and it contains a nonsingular matrix (which can be efficiently found even deterministically), Lemma 21 gives a reduction to finding composition series for matrix algebras which is further reducible to factorization of polynomials over $\mathbb{F}$. It would also be interesting finding maximum rank matrices over very small fields in the rank-one spanned case. The algorithm of [22] does the job when rank-one generators are at hand.

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" (MOE2012-T3-1-009). 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), 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] Alexander L. Chistov, Gábor Ivanyos, and Marek Karpinski. Polynomial time algorithms for modules over finite dimensional algebras. In *ISSAC*, pages 68–74, 1997.

[7] Eric Chitambar, Runyao Duan, and Yaoyun Shi. Multipartite-to-bipartite entanglement transformations and polynomial identity testing. *Physical Reveiw A*, 81(5):052310, 2010.

[8] 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.

[9] Jack Edmonds. Systems of distinct representatives and linear algebra. *J. Res. Nat. Bur. Standards Sect. B*, 71:241–245, 1967.

[10] 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.

[11] David Eisenbud and Joe Harris. Vector spaces of matrices of low rank. *Advances in Mathematics*, 70(2):135–155, 1988.

[12] Michael A. Forbes and Amir Shpilka. Quasipolynomial-time identity testing of noncommutative and read-once oblivious algebraic branching programs. In *FOCS*, 2013.

[13] Marc Fortin and Christophe Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. *Séminaire Lotharingien de Combinatoire*, 52:B52f, 2004.

[14] James Geelen and Satoru Iwata. Matroid matching via mixed skew-symmetric matrices. *Combinatorica*, 25(2):187–215, 2005.

[15] James F. Geelen. Maximum rank matrix completion. *Linear Algebra and its Applications*, 288:211–217, 1999.

[16] 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.

[17] Leonid Gurvits. Quantum matching theory (with new complexity theoretic, combinatorial and topological insights on the nature of the quantum entanglement), 2002.

[18] Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. Syst. Sci.*, 69(3):448–484, 2004.

[19] 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.

[20] 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.

[21] Gábor Ivanyos, Marek Karpinski, Youming Qiao, and Miklos Santha. Generalized wong sequences and their applications to edmonds' problems. In Ernst W. Mayr and Natacha Portier, editors, *STACS*, volume 25 of *LIPIcs*, pages 397–408. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2014.

[22] Gábor Ivanyos, Marek Karpinski, and Nitin Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.

[23] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Computational Complexity*, 13(1-2):1–46, 2004.

[24] Erich Kaltofen. Greatest common divisors of polynomials given by straight-line programs. *J. ACM*, 35(1):231–264, 1988.

[25] László Lovász. On determinants, matchings, and random algorithms. In *FCT*, pages 565–574, 1979.

[26] 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.

[27] Kazuo Murota. *Matrices and matroids for systems analysis*. Springer, 2000.

[28] Richard Rado. A theorem on independence relations. *The Quarterly Journal of Mathematics, Oxford Ser.*, 13(1):83–89, 1942.

[29] 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.

[30] 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.

[31] Leslie G. Valiant. Completeness classes in algebra. In *STOC*, pages 249–261, 1979.

[32] D. J. A. Welsh. On matroid theorems of Edmonds and Rado. *Journal of the London Mathematical Society*, 2(2):251–256, 1970.

[33] Kai-Tak Wong. The eigenvalue problem $\lambda Tx + Sx$. *Journal of Differential Equations*, 16(2):270 – 280, 1974.

[34] 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.