Constructive noncommutative rank computation in deterministic polynomial time over fields of arbitrary characteristics

An addendum to the paper entitled

Non-commutative Edmonds' problem and matrix semi-invariants

Gábor Ivanyos* Youming Qiao† K. V. Subrahmanyam ‡

January 20, 2016

Abstract

We extend our techniques developed in our manuscript mentioned in the subtitle to obtain a deterministic polynomial time algorithm for computing the non-commutative rank together with certificates of linear spaces of matrices over sufficiently large base fields.

The key new idea is a reduction procedure that keeps the blow-up parameter small, and there are two methods to implement this idea: the first one is a greedy argument that removes certain rows and columns, and the second one is an efficient algorithmic version of a result of Derksen and Makam. Both methods rely crucially on the regularity lemma in the aforementioned manuscript, and in this note we improve that lemma by removing a coprime condition there.

1 Introduction: blow-ups and the non-commutative rank

Let $\mathbb{F}$ be a field, and $M(n,\mathbb{F})$ be the linear space of $n\times n$ matrices over $\mathbb{F}$. We call a linear subspace of $M(n,\mathbb{F})$ a *matrix space*. Then, given a matrix space $\mathcal{B}\leq M(n,\mathbb{F})$, its (k,ℓ)-blow-up $\mathcal{B}^{\{k,\ell\}}$ is defined as the matrix space $\mathcal{B}\otimes M(k\times\ell,\mathbb{F})$ in $M(nk\times n\ell,\mathbb{F})$. Although the most relevant blow-ups in this context are square (e.g, of the form $\mathcal{B}^{\{k,k\}}$), non-square blow-ups turned out to be crucial in the reduction techniques in [DM15].

From earlier results, we can define the non-commutative rank of $\mathcal{B}$ to be the maximum over d of $\frac{1}{d}$ times the maximum rank of a matrix from the blow-up $\mathcal{B}^{\{d,d\}}$. An important question is to determine bounds on the blow-up parameter d (as a function of n) which achieves the desired maximum. From the work of [Der01] it is known that $d\leq n^2 2^{O(n^2)}$, over algebraically closed fields of characteristic zero.[1] In [IQS15] it was shown that over large fields of arbitrary characteristic $d\leq n^4 2^{O(n\log n)}$. This bound even holds over fields of characteristic zero.

These exponential bounds turned out to be sufficient in [GGOW15] to compute the non-commutative rank in deterministic polynomial time, over fields of characteristic zero via a previous

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

†Centre for Quantum Computation and Intelligent Systems, University of Technology, Sydney, Australia. (jimmyqiao86@gmail.com)

‡Chennai Mathematical Institute, Chennai, India. (kv@cmi.ac.in).

[1] Derksen's result applies to a wide class of invariant rings.

algorithm of Gurvits [Gur04]. In [GGOW15] the following open problems were posed (1) a polynomial time algorithm for the problem over finite fields, and (2) a search version of the problem - explicitly exhibiting a matrix of rank rd in the d-th blow-up (even over fields of characteristic 0) and a proof that the non commutative rank is at most r. Very recently, Derksen and Makam proved in [DM15] that it suffices to take the maximum over d between 1 and $n-1$ (in arbitrary characteristic).

In this note we combine methods from [DM15] with those from [IQS15] to obtain an efficient constructive version of the former result. After we wrote up the first version of this note we discovered that a very simple observation gives us the result, without having to use the results from Derksen and Makam. This argument also gives a different proof that the nullcone is defined by polynomials with degree a polynomial in n.

We say that a subspace V of $\mathbb{F}^n$ is an s-shrunk subspace of $\mathcal{B}$ if $\dim \mathcal{B}V \leq \dim V - s$. If $\mathcal{B}$ has an s-shrunk subspace, then no matrix from $\mathcal{B}$ can have rank larger than $n-s$. As by a square d-blow-up, s-shrunk subspaces are blown up to ds-shrunk subspaces, the existence of an s-shrunk subspace implies an upper bound $n-s$ for the noncommutative rank of $\mathcal{B}$. Fortin and Reutenauer showed in [FR04] that the minimum of the number $n-s$ such that there exist s-shrunk subspaces coincides with the non-commutative rank.

Our main algorithmic result will construct objects witnessing the non-commutative rank from both sides.

Theorem 1. *Let $\mathcal{B} \leq M(n, \mathbb{F})$ be a matrix space given by a basis, and suppose $|\mathbb{F}| = n^{\Omega(1)}$. Suppose that $\mathcal{B}$ has (a priori unknown) non-commutative rank r. Then there is a deterministic algorithm using $n^{O(1)}$ arithmetic operations over $\mathbb{F}$ that constructs a matrix of rank rd in a blow-up $\mathcal{B}^{\{d,d\}}$ for some $d \leq r+1$ as well as an $(n-r)$-shrunk subspace of $\mathbb{F}^n$ for $\mathcal{B}$. When $\mathbb{F} = \mathbb{Q}$, the final data as well as all the intermediate data have size polynomial in the size of the input data and hence the algorithm runs in polynomial time.*

Remark 2.

- For the above theorem, we use our method in Section 4.1. If the method of Derksen and Makam [DM15] is used as in Section 4.2, then in the above theorem we can improve the parameter slightly to $d \leq r-1$ instead of $d \leq r+1$.
- Polynomiality of the algorithm can also be proved for a wide range "concrete" base fields $\mathbb{F}$. These include sufficiently large finite fields, and also number fields and transcendental extensions of constant degree over finite fields.

In particular, the non-commutative rank can be computed in deterministic polynomial time in positive characteristic as well, assuming that the ground field is sufficiently large. This then settles the aforementioned two open problems from [GGOW15].

Our result also settles a question of Gurvits [Gur04], asking if it is possible to decide, efficiently, in fields of positive characteristic, whether or not there exists a nonsingular matrix in a matrix space having the so called Edmonds–Rado property. A matrix space has the Edmonds-Rado property if it satisfies the promise that it either contains a nonsingular matrix, or it shrinks some subspace. Yet another way to view this result is that we settle the determinant identity testing for matrix spaces satisfying the Edmonds Rado property, when the field size is as stated in the hypothesis.

However we can say more, even when the base field is a "too small" finite field. We can then apply Theorem 1 to a sufficiently large extension. We can embed the extension field as a maximal commutative subfield of the full matrix algebra of an appropriate degree over the original field, using the regular representation. We obtain the following.

Corollary 3. *Let $\mathcal{B} \leq M(n, \mathbb{F})$ be a matrix space given by a basis, where $\mathbb{F}$ is a finite field. Suppose that $\mathcal{B}$ has (a priori unknown) non-commutative rank r. Then there is a deterministic polynomial time algorithm that constructs a matrix of rank rd in a blow-up $\mathcal{B}^{\{d,d\}}$ for some $d \leq rO(\log_{|\mathbb{F}|} r)$ as well as an $(n-r)$-shrunk subspace of $\mathbb{F}^n$ for $\mathcal{B}$.*

So we settle Gurvits [Gur04] question completely, and also solve the determinant identity testing for matrix spaces satisfying the Edmonds-Rado property.

2 Regularity and incrementing rank via further blow-up

The following lemma (Lemma 11 from [IQS15]) states that in a (square) blow-up it can be efficiently arranged that the structure of a (seemingly) maximum rank matrix shows some minimal consistencies with the block structure of the blow-up. Namely, it can be enforced that there is a full rank sub-matrix which is a blow-up from the space of sub-matrices from the original matrix space obtained by restricting the matrices to a "window" determined by a set of rows and set of columns. In particular, the rank in the blow-up can be efficiently "rounded" up to the next multiple of the blowing parameter. It is natural to scale down this rank with the blowing parameter and to consider the result as an approximation of the non-commutative rank of the original space.

Lemma 4 (Regularity of blow-ups, [IQS15, Lemma 11])**.** *For $\mathcal{B} \leq M(n, \mathbb{F})$ and $\mathcal{A} = \mathcal{B}^{\{d,d\}}$, assume that $|\mathbb{F}| > rd$ (which is obviously true when $\mathrm{char}(\mathbb{F}) = 0$). Given a matrix $A \in \mathcal{A}$ with $\mathrm{rk}A > (r-1)d$, there exists a deterministic algorithm that returns $\widetilde{A} \in \mathcal{A}$ and an $r \times r$ window W in $\widetilde{A}$ s.t. W is nonsingular (of rank rd). This algorithm uses $\mathrm{poly}(nd)$ arithmetic operations and, over $\mathbb{Q}$, the algorithm runs in polynomial time. In particular, all intermediate numbers have bit lengths polynomial in the input size.*

Here by an $r \times r$ window of a matrix from the blow-up we mean a $dr \times dr$ sub-matrix whose columns and rows are obtained by blowing up r rows and r columns from the $n \times n$ setting. To be more precise, suppose $\mathbf{i} = (i_1, \ldots, i_r)$, $\mathbf{j} = (j_1, \ldots, j_r)$ are two sequences of integers, where $1 \leq i_1 \leq \cdots \leq i_r \leq n$ and $1 \leq j_1 \leq \cdots \leq j_r \leq n$. For a matrix $A \in M(n, \mathbb{F}) \otimes M(d, \mathbb{F})$, the $r \times r$ *window* indexed by $\mathbf{i}, \mathbf{j}$ is the sub-matrix of A consisting of the blocks indexed by (i_k, j_ℓ), $k, \ell \in [r]$, the $r \times r$ *window* indexed by $\mathbf{i}, \mathbf{j}$.

The cases 1) $\mathrm{char}(\mathbb{F}) = 0$, 2) $\mathrm{char}(\mathbb{F})$ and d are coprime, and $|\mathbb{F}| > 2rd$ were settled in [IQS15, Lemma 11]. The main issue with the case when d is divisible by $\mathrm{char}(\mathbb{F})$ was that then the proof would have required an efficient construction of an appropriate Artin–Schreier–Witt extension of $\mathbb{F}_p(x)$, which was unavailable to us at the time of writing [IQS15]. Now we have such a construction. So we include a complete proof of Lemma 4 in this note. We borrow the idea from the cases solved in the earlier paper and, to cover the remaining case, we prove Lemma 7 which gives an efficient construction of an Artin–Schreier–Witt extension of $\mathbb{F}_p(x)$.

The main technical ingredient of our algorithm will be the following result from [IQS15] (Theorem 13 of ibid.). It states that either a shrunk subspace witnessing that the (scaled down) rank of a matrix in a blow-up reaches the non-commutative rank or a matrix in a larger blow-up having larger scaled down rank can be efficiently constructed.

Theorem 5 ([IQS15, Theorem 13])**.** *Let $\mathcal{B} \leq M(n, \mathbb{F})$ and let $\mathcal{A} = \mathcal{B}^{\{d,d\}}$. Assume that we are given a matrix $A \in \mathcal{A}$ with $\mathrm{rk}(A) = rd$, and $|\mathbb{F}|$ is $\Omega(ndd')$, where $d' > r$ is any positive integer. There exists a deterministic algorithm that returns either an $(n-r)d$-shrunk subspace for*

$\mathcal{A}$ (equivalently, an $(n-r)$-shrunk subspace for $\mathcal{B}$), or a matrix $B \in \mathcal{A} \otimes M(d', \mathbb{F})$ of rank at least $(r+1)dd'$. To be more specific, in this case an $(r+1) \times (r+1)$ window is also found such that the corresponding $(r+1)dd'$ times $(r+1)dd'$ sub-matrix of B has full rank. This algorithm uses $\mathrm{poly}(ndd')$ *arithmetic operations and, over $\mathbb{Q}$, all intermediate numbers have bit lengths polynomial in the input size.*

The sentence on the $(r+1) \times (r+1)$ window is not explicitly stated in [IQS15]. However, the algorithm in its proof contains, as a last step, a call to the method behind Lemma 4. Also, the theorem was stated only under the assumption that d was not divisible by $\mathrm{char}(\mathbb{F})$ because of this last call. As the algorithm up to this step constructs a matrix of rank greater than rdd', the complete version of Lemma 4 makes it possible to dispense with that assumption.

Finding an sd-shrunk subspace for the $\mathcal{B}^{\{d,d\}}$ is equivalent to finding an s-shrunk subspace for $\mathcal{B}$ because of the following simple observations. Firstly, for every s-shrunk subspace U of $\mathbb{F}^n$ the subspace $U \otimes \mathbb{F}^d$ for $\mathcal{B}$ is obviously an sd-shrunk subspace for $\mathcal{B}^{\{d,d\}}$. Conversely, a s'-shrunk subspace for $\mathcal{B}^{\{d,d\}}$ can be easily embedded into a subspace of the form $U \otimes \mathbb{F}^d$ where U is an s-shrunk subspace for $\mathcal{B}$ with $sd \geq s'$.

The next lemma concerns a simple procedure originally developed in [dGIR96] for reducing data for an algorithm for a task in a different, although not completely unrelated context. See also the appropriate paragraph in Subsection 4.1 of [IQS15] for a brief description. This technique is extensively used in the above mentioned two results for handling extension fields if necessary. It is also used to keep data small over infinite base fields. Here it can also be applied to keep the entries of the d by d matrices defining the large rank matrix from the d-blow-up in a finite (but sufficiently large) subset of the base fields.

Lemma 6 (Data reduction, [dGIR96,IQS15])**.** *Let $\mathcal{A} \leq M(k \times \ell, \mathbb{F})$ be given by a basis $A_1, \ldots, A_m$. Let S be a subset S of $\mathbb{F}$ of size at least $r+1$. Suppose that we are also given a matrix $B = \sum \beta_i A_i \in \mathcal{A}$ of rank at least r. Then we can find $A = \sum \alpha_i A_i$ of rank also at least r with $\alpha_i \in S$. The algorithm uses* $\mathrm{poly}(k, \ell, r)$ *rank computations for matrices of the form $\sum \gamma_i A_i$ where $\gamma_i \in \{\beta_1, \ldots, \beta_m\} \cup S$.*

3 Regularity of blow-ups

This section is devoted to a proof of Lemma 4.

3.1 Ingredients: cyclic extensions and cyclic division algebras

Recall that a cyclic extension of a field K is a finite Galois extension of K having a cyclic Galois group. By constructing a cyclic extension L we mean constructing the extension as an algebra over K, e.g., by giving an array of *structure constants* with respect to a K-basis for L defining the multiplication on L as well as specifying a generator of the Galois group, e.g, by its matrix with respect to a K-basis. Recall that for a finite dimensional algebra $\mathcal{A}$ over the field K, a common way to specify the multiplication is using an array of structure constants with respect to a K-basis $A_1, \ldots, A_d$. These are d^3 elements γ_{ijk} of K such that $A_i A_j = \sum_{k=1}^{d} \gamma_{ijk} A_k$. Then we can represent elements of $\mathcal{A}$ by the vectors of their coordinates in terms of the basis $A_1, \ldots, A_d$. The size of the data representing the structure constants give some control over the size of data representing the product of element. Perhaps the easiest example is the case where K is the function field $\mathbb{F}_p(Z)$ and the structure constants happen to be polynomials in $\mathbb{F}_p[Z]$. Then for two elements of $\mathcal{A}$ with

their coordinates being polynomials in $\mathbb{F}_p[Z]$, their product will have also polynomial coordinates, and the degrees of the coordinates of the product are upper bounded by the sum of the maximum degrees of coordinates the factors plus the maximum degree of the structure constants.

Lemma 7. *Given a prime p and an integer $s \geq 1$, one can construct in time* $\mathrm{poly}(p^s)$ *a cyclic extension K_s of $\mathbb{F}_p(Z)$ of degree p^s such that $\mathbb{F}_p$ is algebraically closed in K_s. The field K_s will be given in terms of structure constants with respect to a basis over $\mathbb{F}_p(Z)$, and the generator σ for the Galois group will be given by its matrix in terms of the same basis. The structure constants as well as the entries of the matrix for σ will be polynomials in $\mathbb{F}_p[Z]$ of degree* $\mathrm{poly}(p^s)$.

Proof. First we briefly recall the general construction given in Section 6.4 of [Ram54]. This, starting from a field K_0 of characteristic p, recursively builds a tower $K_0 < K_1 < \ldots < K_s$ of fields such that K_j is a cyclic extension of K_0 of degree p^j. Assume that K_s together with a K_0-automorphism σ_s of order p^s has already been constructed. (Initially let σ_0 be the identity map on K_0.) Then for any element $\beta_s \in K_s$ with $\mathrm{Tr}_{K_s:K_0}(\beta_s) = 1$ and for any $\alpha_s \in K_s$ such that $\alpha_s^{\sigma_s} - \alpha_s = \beta_s^p - \beta_s$ the polynomial $X^p - X - \alpha_s$ is irreducible in $K_s[X]$. (Existence of α_s with the required property follows from the additive Hilbert 90.) Put $K_{s+1} = K_s[X]/(X^p - X - \alpha_s)$ and let $\omega_{s+1} \in K_{s+1}$ be the image of of X under the projection $K_s[X] \to K_{s+1}$. Then σ_s extends to a K_0-automorphism σ_{s+1} of degree p^{s+1} of K_{s+1} such that $\omega_{s+1}^{\sigma_{s+1}} = \omega_{s+1} + \beta_s$. This gives a cyclic extension of degree p^{s+1}.

Now we specify some details of a polynomial time construction for $K_0 = \mathbb{F}_p(Z)$ following the method outlined above. In the first step we take $\beta_0 = 1$, and, in order to guarantee that the only elements in K_1 which are algebraic over $\mathbb{F}_p$ is F_p (we also use the phrase F_p is algebraically closed in K_1 when this property holds), we take $\alpha_0 = Z$. Then K_1 is a pure transcendental extension of $\mathbb{F}_p$. As K_s/K_0 is a cyclic extension of oder p^s, it has a unique subfield which is an order p extension of K_0. This must be K_1. Then $\mathbb{F}_p$ has no proper finite extension in K_s as otherwise K_0 would also have another degree p extension.

We consider the following K_0-basis for K_s:

$$\Gamma_s = \left\{ \prod_{j=1}^{s} \omega_j^k, \quad (k = 0, \ldots, p-1) \right\},$$

where ω_j is a root of $X^p - X - \alpha_{j-1}$ in K_j. We claim that $\mathrm{Tr}_{K_j:K_{j-1}}(\omega_j^{p-1}) = -1$. Indeed, in the K_{j-1}-basis $\omega_j^0, \ldots, \omega_j^{p-1}$ for K_j, in the matrix of multiplication by ω_j^{p-1} the diagonal entries consist of $p-1$ ones and one zero. Therefore $\mathrm{Tr}_{K_j:K_{j-1}}(\omega_j^{p-1}\gamma) = -\gamma$ for every $\gamma \in K_{j-1}$, whence $\mathrm{Tr}_{K_j:K_0}(\omega_j^{p-1}\gamma) = -\mathrm{Tr}_{K_{j-1}:K_0}(\gamma)$. Now by induction we obtain $\mathrm{Tr}_{K_j:K_0} \prod_{i=1}^{j} \omega_i^{p-1} = (-1)^j$. Therefore in each step (when $j > 0$) we can choose $\beta_j = (-1)^j \prod_{i=1}^{j} \omega_i^{p-1}$ and α_j thereafter, following the construction in the standard proof of the additive Hilbert 90. Specifically, we set

$$\alpha_j = (-1)^{j+1} \sum_{k=1}^{p^j-1} \beta_j^{\sigma_j^k} \left(\sum_{\ell=0}^{k-1} (\beta_j^p - \beta_j)^{\sigma_j^\ell} \right). \tag{1}$$

Then $\alpha_j^{\sigma_j} - \alpha_j = \beta_j^p - \beta_j$. Notice that α_j is a sum of terms with each of which, up to a sign, is a product of at most $p+1$ conjugates $\beta_j^{\sigma_j^\ell}$ (with various ℓs) of β_j ($\ell \leq p^j$)

Assume by induction that the structure constants of K_j with respect to the basis Γ_j are polynomials from $\mathbb{F}_p[Z]$ of degree at most Δ_j and the same holds for the entries of the matrix of σ_j^ℓ for every $1 \leq \ell < p^j$ (written in the same basis). For $j = 1$ this holds with $\Delta_1 = 1$. (To see this, observe that for $0 \leq k, \ell < p$, the product $\omega_1^k \omega_1^\ell$ is the basis element of $\omega_1^{k+\ell}$ if $k + \ell < p$, while otherwise it equals the sum $\omega_1^{k+\ell-p+1} + Z\omega_1^{k+\ell-p}$.) Then, if we express α_j in terms of the basis Γ_j using Eq. 1, we obtain that its coordinates are polynomials of degree at most $(2p+1)\Delta_j$. This is because $(-1)^j \beta_j \in \Gamma_j$, whence $\beta_j^{\sigma^\ell}$ has coordinates of polynomials of degree bounded by Δ_j. In Eq. 1, we have the products of at most $p+1$ such elements, so the result will have polynomial coordinates of degree at most $(2p+1)\Delta_j$.

Now consider the product of two elements $\omega_{j+1}^k \gamma_1$ and $\omega_{j+1}^\ell \gamma_2$ of Γ_{j+1}. Here $k, \ell < p$ and $\gamma_1, \gamma_2 \in \Gamma_j$. The coordinates of the product $\gamma_1\gamma_2$ with respect to Γ_j are polynomials of degree at most Δ_j. The same holds for the product $\omega_{j+1}^{k+\ell}\gamma_1\gamma_2$ if $k + \ell < p$. If $k + \ell > p$, then $\omega_{j+1}^{k+\ell} = \omega_{j+1}^p \omega_{j+1}^{k+\ell-p} = (\omega_{j+1} + \alpha_j)\omega_{j+1}^{k+\ell-p}$, whence $\omega_{j+1}^{k+\ell}\gamma_1\gamma_2$ is the sum of $\omega_{j+1}^{1+k+\ell-p}\gamma_1\gamma_2$ and $\alpha_j\gamma_1\gamma_2$. The former term has coordinates of degree at most Δ_j, the coordinates of the latter are polynomials of degree at most $(2p+1)\Delta_j + \Delta_j + \Delta_j = (2p+3)\Delta_j$.

Now consider the conjugate of $\omega_{j+1}^k \gamma$ by σ_{j+1}^ℓ, where $1 \leq \ell < p^{j+1}$, $1 \leq k \leq p-1$ and $\gamma \in \Gamma_j$. This conjugate is $(\omega_{j+1}^{\sigma_{j+1}^\ell})^k \gamma^{\sigma_{j+1}^\ell}$. The second term equals $\gamma^{\sigma_j^\ell}$ which has coordinates of degree at most Δ_j. To investigate the first term, recall that $\omega_{j+1}^{\sigma_{j+1}} = \omega_{j+1} + \beta_j$, whence

$$\omega_{j+1}^{\sigma_{j+1}^\ell} = \omega_{j+1} + \sum_{r=0}^{\ell-1} \beta_j^{\sigma_j^r}$$

The element $\delta = \sum_{r=0}^{\ell-1} \beta_j^{\sigma_j^r}$, expressed in terms of Γ_j, has again polynomial coordinates of degree at most Δ_j. Then $(\omega_{j+1}^{\sigma_{j+1}^\ell})^k$ is the sum (with binomial coefficients) of terms of the form $\omega_{j+1}^r \delta^{k-r}$. The power δ^{k-r} has coordinates of degree at most $(k-r)\Delta_j + (k-r-1)\Delta_j \leq (2p-1)\Delta_j$ in terms of Γ_j, whence we conclude that $(\omega_{j+1}^{\sigma_{j+1}^\ell})^k$ has, in terms of Γ_{j+1} polynomial coordinates of degree at most $(2p-1)\Delta_j$. It follows that the matrix of any power of σ_{j+1} has polynomial entries of degree at most $2p\Delta_j$.

We obtained that the function $(2p+3)^s = \mathrm{poly}(p^s)$ is an upper bound for both the structure constants and for the matrices of the powers of σ_s. □

Lemma 8. *Let p be a prime such that $d = d_1 p^s$ where d_1 is the largest divisor of d not divisible by p. In particular p is an arbitrary prime and $s = 0$ if the characteristic of K is zero while otherwise p is the characteristic of K. Assume that K contains a known d_1th root of unity ζ. Then a cyclic extension L degree d of $K(X_1, X_2)$ can be computed using* $\mathrm{poly}(d)$ *arithmetic operations. L will be given by structure constants with respect to a basis, and the matrix for a generator of the Galois group in terms of the same basis will also be given. All the output entries (the structure constants as well as the entries of the matrix representing the Galois group generator) will be polynomials of degree* $\mathrm{poly}(d)$ *in $K[X_1, X_2]$. Furthermore for $K = \mathbb{Q}[\sqrt[d_1]{1}]$, the bit complexity of the algorithm (as well as the size of the output) is* $\mathrm{poly}(d)$.

Proof. Put $L_1 = K(Y)$ and $X_1 = Y^{d_1}$. Then $1, Y, \ldots, Y^{d_1}$ are a $K(X_1)$-basis for L_1 with $Y^i Y^j = Y^{i+j}$ if $i + j \leq d_1$ and $X_1 Y^{i+j-d_1}$ otherwise. Also, the linear extension σ_1 of the map sending Y^j to $\zeta^j Y^j$ is an automorphism of degree d_1. Let $d_2 = p^s$. For a cyclic extension L_2 of $K(X_2)$

(if $d_2 > 1$) we compute the prime field $\mathbb{F}_p$ of K by testing the multiples of the identity element, and then take the tensor product of K with the cyclic extension of degree d_2 of $\mathbb{F}_p(X_2)$ obtained by the construction of Lemma 7 over $\mathbb{F}_p$. We also obtain the matrix a generator σ_2 of the Galois group. Then put $L = L_1 \otimes_K L_2$. It contains a copy of $K(X_1, X_2) \cong K(X_1) \otimes_K K(X_2)$. We take the product basis for the structure constants and for matrix representation of the automorphism $\sigma_1 \otimes \sigma_2$. □

The following statement follows from Wedderburn's theorem characterizing cyclic division algebras (see e.g. [Lam91, Theorem (14.9)]) as shown on Page 240 of [Lam91].

Fact 9. *Let L be a cyclic extension of degree d of the field K'. Let σ be a generator of the Galois group. We consider the transcendental extension $L(Z)$ of L. Then σ extends to an automorphism (denoted again by σ) of $L(Z)$ such that the fixed field of σ is $K'(Z)$. Thus $L(Z)$ is a cyclic extension of $K'(Z)$. Consider the $K'(Z)$-algebra D generated by (a basis for) L and by an element U with relations $U^d = Z$ and $Ua = a^\sigma U$ ($\forall a \in L(Z)$, or, equivalently $\forall a \in$ the basis for L). Then D is a central division algebra of index d over $K(Z)$.*

Lemma 10. *Let K and L be as in Lemma 8. Then one can construct a $K(X_1, X_2, Y)$-basis Γ of $M(n, K(X_1, X_2, Y))$ such that the $K(X_1, X_2, Y^d)$-linear span of Γ is a central division algebra over $K(X_1, X_2, Y^d)$ of index d using* poly(d) *arithmetic operations in K. Furthermore for $K = \mathbb{Q}[\sqrt[d_1]{1}]$, the bit complexity of the algorithm (as well as the size of the output) is also* poly(d).

Proof. Let $K' = K(X_1, X_2)$, put $Z = Y^d$ and let D be a central division algebra over K' as in Fact 9. We also consider the commutative subfield $L(Z)$ of D which is a cyclic extension of degree d of $K'(Z)$. Existence of a $K(Z)$-subalgebra D' of $M(d, K'(Y))$ isomorphic to D follows, e.g., from Theorem (14.7) of [Lam91]. To construct a basis Γ for such a matrix algebra D' efficiently, consider the $K(Z)$-basis for D consisting of the products $A_i U^j$ $(i, j = 1, \ldots, d)$. (Here we assume that $A_1, \ldots, A_d$ correspond in $L(Z)$ to the basis elements of L for which Lemma 8 constructed the structure constants.) This is also a $K(Y)$-basis for the algebra $D'' = K(Y) \otimes D$. Consider also the element $U_0 = \frac{1}{Y}U$. Then $U_0^d = 1$. As the elements U_0^j are linearly independent over $K(Y)$ and hence over $K(Z)$ as well, we have that $E = U_0 + U_0^2 + ... + U_0^{d-1} + U_0^d$ is nonzero. As $U^j E = Y^j E$, we have $A_i E$ $(i = 1, \ldots, d)$ are a $K(Y)$-basis for the left ideal $D''E$ of dimension d. Now the action of D'' on this left ideal gives a matrix representation for D''. Let γ_{kij} be the structure constants for the multiplication of L (and of $L(Z)$):

$$A_k A_i = \sum_{j=1}^{d} \gamma_{kij} A_j \quad (k, i = 1, \ldots, d).$$

Also, let $\delta_{\ell ij}$ be the entries of the matrix of the ℓth power of the generator σ of the Galois group:

$$A_i^{\sigma^\ell} = \sum_{j=1}^{d} \delta_{\ell ij} A_j \quad (\ell, i = 1, \ldots, d).$$

(Notice that the matrix $(\delta_{\ell ij})_{ij}$ is the ℓth power of $(\delta_{1ij})_{ij}$, whence its degree is also bounded by poly(d).) Then

$$A_k A_i E = \sum_{j=1}^{d} \gamma_{kij} A_j E$$

and

$$U^\ell A_i E = A_i^{\sigma^\ell} U^\ell E = A_i^{\sigma^\ell} Y^\ell E = Y^\ell \sum_{j=1}^{d} \delta_{\ell ij} A_j E.$$

Thus the matrix of the action of A_k has entries γ_{kij} and the matrix of the action of U^ℓ has entries $Y^\ell \delta_{\ell ij}$. Then the action of of $U^\ell A^k$ can be obtained as the product of these two matrices. Let Γ consist of all such d^2 products, and the proof is concluded. □

3.2 Proof of regularity

We shall make use of Lemma 10 as follows. First, we take a (transcendental) extension $\mathbb{K}$ of $\mathbb{F}$ so that we can construct a central division algebra D of index d over $\mathbb{K}(Z) = \mathbb{K}(Y^d)$ by Lemma 8 and Fact 9. Then, instead of the blow-up $\mathcal{B}^{\{d,d\}} = \mathcal{B} \otimes M(d, \mathbb{F})$ we consider the blow-up $\mathcal{B} \otimes M(d, \mathbb{K}(Y))$. Let D' be a $\mathbb{K}(Z)$-subalgebra of $M(d, \mathbb{K}(Y))$ such that $\mathbb{K}(Y) \otimes D' = M(d, \mathbb{K}(Y))$ and that D' is a division algebra. Then both $\mathcal{B}^{\{d,d\}}$ and $\mathcal{B} \otimes_{\mathbb{F}} D'$ span, as a $\mathbb{K}(Y)$-linear space, the blow-up $\mathcal{B} \otimes_{\mathbb{F}} M_d(\mathbb{K}(Y))$.

Claim 11. *Let $\mathbb{K}$ and $D' \subset M(d, \mathbb{K}(Y))$ be as above. Then every matrix in $M(n, \mathbb{F}) \otimes D' \subset M(d, \mathbb{K}(Y))$ has rank (as a matrix over $\mathbb{K}(Y)$) divisible by d.*

Proof. Consider the $\mathbb{K}$-linear space $\mathbb{K}(Y)^{dn} \cong \mathbb{K}(Z)^{d^2 n}$ over $\mathrm{M}(n, \mathbb{F}) \otimes_{\mathbb{F}} D' = M(n, \mathbb{K}(Z)) \otimes_{\mathbb{K}(Z)} D'$. Since $D' \otimes_{\mathbb{K}(Z)} D'^{op} \cong M(d^2, \mathbb{K}(Z))$ [Lam91, Corollary (15.5)], it follows that the centralizer of this action is the opposite algebra D'^{op}. Therefore the image of $A'\mathbb{K}^{dn}$ of any $A' \in M(n, \mathbb{F}) \otimes_{\mathbb{F}} D'$ is a D'^{op}-submodule, whence its dimension over $\mathbb{K}(Z)$ is divisible by d^2. It follows that the dimension over $\mathbb{K}(Y)$ is divisible by d. □

The claim enables us to "round up" ranks of an already constructed matrix in the blow-up to the next multiple of d. We outline a method for this task below.

Let A be a matrix from $\mathcal{B}^{\{d,d\}}$ of rank greater than $(r-1)d$ where $r \leq n$. Then A, considered as a matrix over $\mathbb{K}(Y)$ has also rank greater than $(r-1)d$. Now consider a basis for D' over $\mathbb{K}(Z)$. Then A can be expressed as a linear combination (with coefficients over $\mathbb{K}(Y)$) of these basis elements. We use the method of Lemma 6 to find coefficients from $\mathbb{K}(Z)$ (or even from $\mathbb{F}$) such that the combination A' of the basis element for D' has rank also larger than $(r-1)d$. We have $A' \in \mathcal{B} \otimes D'$, whence by Claim 11, the rank of A' is at least rd. Then we express A' as a linear combination of elements – with coefficients from $\mathbb{K}(Y)$ – of an $\mathbb{F}$-basis of $\mathcal{B}^{\{d,d\}}$ which is also a $\mathbb{K}(Y)$-basis for $\mathcal{B} \otimes M(d, \mathbb{K}(Y))$. Then we use again the algorithm of Lemma 6 to replace these coefficients to elements of $\mathbb{F}$ to find a matrix $A'' \in \mathcal{B}$ of rank at least rd.

An efficient realization of this method gives the following.

Lemma 12. *Let $\mathcal{B} \leq M(n, \mathbb{F})$ and assume that we are given $A \in \mathcal{A} = \mathcal{B}^{\{d,d\}}$ of rank larger than $(n-1)d$. Then, provided that the size of $\mathbb{F}$ is $(nd)^{\Omega(1)}$, using $\mathrm{poly}(nd)$ operations in $\mathbb{F}$ we can find a matrix of rank nd in $\mathcal{A}$. Furthermore, if $\mathbb{F} = \mathbb{Q}$ then the bit complexity of finding such a matrix is polynomial in the size of the input data (these are the entries of a basis $B_1, \ldots, B_m$ for $\mathcal{B}$ and the entries of the matrices $C_i \in M(d, \mathbb{F})$ such that $A = \sum_{i=1}^{m} A_i \otimes C_i$).*

Proof. We describe the details of the critical ingredients for turning the method described above into an efficient algorithm.

To begin with, we need an extension $\mathbb{K}$ of our base field $\mathbb{F}$ such that we can construct a central division algebra over $\mathbb{K}$ as in Fact 9. For this a finite extension $\mathbb{F}'$ of $\mathbb{F}$ containing a (known) primitive d_1th root of unity ζ (here d_1 is char($\mathbb{F}$)-free part of d) and $\mathbb{K} = \mathbb{F}'(X_1, X_2)$ would be sufficient.

Then $\mathbb{K}(Y) = \mathbb{F}'(X_1, X_2, Y)$ is a pure transcendental extension of $\mathbb{F}'$. Then computation of the rank of $nd \times nd$ matrices with entries from $\mathbb{K}(Y)$ can be accomplished as follows. Assume that the entries of the given matrix are represented by quotients of polynomial form $\mathbb{F}'[X_1, X_2, Y]$. Then the size of the input is the total size of these $2n^2d^2$ polynomials. We multiply all the entries by an easily computable common multiple (e.g., the product) of their denominators to obtain a matrix with polynomial entries from $\mathbb{F}'[X_1, X_2, Y]$. The data describing this matrix have size polynomial in the size of the input data. In particular, the degree of the determinant of any sub-matrix is bounded by a polynomial s in nd and the sum of the degrees appearing in the input entries. We assume that $\mathbb{F}'$ (and even $\mathbb{F}$) has more than s element. Then from $(s+1)^3$ specializations by elements of subset of size $s+1$ of $\mathbb{F}'$ at least one gives a matrix with entries $\mathbb{F}'$ having the same rank as the original matrix. Thus to compute the rank over $\mathbb{K}(Y)$ can be accomplished by computing the rank over $\mathbb{F}'$ of polynomially many specializations.

The bounds given in Lemma 10 guarantee that all the degrees of polynomials we encounter are bounded by some polynomial of nd and, over the rationals the bit size of the coefficients of these polynomials also remain bounded in the size of the data describing the input basis for $\mathcal{B}$.

There is one issue regarding the construction of Kummer extensions. Namely, constructing $\mathbb{F}' = \mathbb{F}[\zeta]$ would require factoring the polynomial $x^{d_1} - 1$ over $\mathbb{F}$, a task which cannot be accomplished using basic arithmetic operations. To see that this is indeed an issue notice that a black-box field may contain certain "hidden" parts of cyclotomic fields. (Of course, over certain concrete fields, such as the rationals, number fields or finite fields of small characteristic, this can be done in polynomial time. However, even over finite fields of large characteristic no deterministic polynomial time solution to this task is known at present.)

To get around this issue, one can perform the required computations over an appropriate ideal R of the algebra $C = \mathbb{F}[x]/(x^{d_1} - 1)$ in place $\mathbb{F}'$ as if R were a field. To be specific, as d_1 is not divisible by the characteristic, we know that C is semisimple – actually it is isomorphic to a direct sum of ideals each of which is isomorphic to the splitting field $\mathbb{F}[\sqrt[e]{1}]$ of the polynomial $x^e - 1$ for some divisor e of d_1 and the pojection of x to such an ideal is a primitive eth root of unity. It follows that if we compute the ideal J of C generated by the annihilators of $x^e - 1$ for all proper divisor e of d_1 then $R = C/J$ is isomorphic to the direct sum of copies of the splitting field $\mathbb{F}'$ of $x^{d_1} - 1$ and the projection of x to each component is a primitive d_1th root of unity. And this property is inherited by any proper factor of R. The procedure (whose critical part is the computation of the rank of specialized matrices via, e.g., Gaussian elimination) using R instead of $\mathbb{F}'$ fail if and only if a zero divisor in R is found. In that case we replace R with the factor of R by its ideal generated by the zero divisor and restart the computation.

Note that a similar issue, namely that a black box field may even contain infinite algebraic extensions of its subfields has been circumvented by using the transcendental extension $\mathbb{K} = \mathbb{F}'(X, Y)$ in the construction of cyclic extensions (Lemma 8). □

We are now in a position to finish the proof of Lemma 4.

Proof of Lemma 4. If $\mathbb{F}$ is a finite field of size poly(nd), if necessary we replace $\mathbb{F}$ with an extension of size still poly(nd) but large enough for Lemma 12.

The proof goes by induction on r. To see the initial case $r = 1$, let B be any nonzero matrix from $\mathcal{B}$. Assume that the (i, j)th entry of B is nonzero. Then the (i, j)th block of $B \otimes I$ is a nonzero $d \times d$ scalar matrix.

For the inductive step, assume $r > 1$. By the induction hypothesis, we can find a matrix $A' \in \mathcal{A}$ and an $(r-1) \times (r-1)$ nonsingular window in A'. We assume w.l.o.g. that the window corresponds to row and column indices $1, \ldots, r-1$, that is, the nonsingular sub-matrix of A' consists of its upper left $(r-1)^2$ blocks. It is easy to see that if S is a subset of $\mathbb{F}$ of size at least $2rd + 1$ then for some $(\lambda, \mu) \in S \times S$, we have that the upper left $(r-1)d \times (r-1)d$ block of $\lambda A' + \mu A$ is still nonsingular and at the same time $\lambda A' + \mu A$ has rank larger than $(r-1)d$. We replace A with such a $\lambda A' + \mu A$. We can again w.l.o.g. assume that already the upper left r^2 blocks of A form a sub-matrix of rank larger that $(r-1)d$. We apply the algorithm for Lemma 12 to this upper left $rd \times rd$ sub-matrix of A and to the corresponding $r \times r$ window to obtain a matrix of rank at least rd.

As a final step, if we have extended our base field than we can go back to the original field using the method of Lemma 6. □

4 Blow-up reduction tools

To obtain the algorithm for Theorem 1, the regularity lemma needs to be accompanied with a reduction procedure that keeps the blow-up parameter small. One such procedure can be obtained by making effective the method of Derksen and Makam, who were the first to observe that it suffices to consider blow-ups of size at most $n-1$. However there is a much simpler procedure, at the small price of considering blow-ups of size at most $n+1$ instead of $n-1$. We give that proof first, and later we also present a constructive version of the Derksen and Makam result.

4.1 A greedy argument

Lemma 13. *Let $\mathcal{B} \leq M(n, \mathbb{F})$, and $d > n + 1$. Assume we are given a matrix $A \in \mathcal{B}^{\{d,d\}}$ of rank dn. Then there exists a deterministic polynomial-time procedure that constructs $A' \in \mathcal{B}^{\{d-1,d-1\}}$ of rank $(d-1)n$.*

Proof. Let A'' be an appropriate $(d-1)n \times (d-1)n$ submatrix of A corresponding a matrix in $\mathcal{B}^{\{d-1,d-1\}}$. We claim A'' is of rank $> (d-1)(n-1)$. Suppose not, as A is obtained from A'' from adding n rows and then n columns, and $d > n+1$, we have $\mathrm{rk}(A) \leq \mathrm{rk}(A'') + 2n \leq dn - d - n + 1 + 2n < dn$, a contradiction. Now that $\mathrm{rk}(A'') > (d-1)(n-1)$, using Lemma 4, we obtain $A' \leq \mathcal{B}^{\{d-1,d-1\}}$ of rank $(d-1)n$. □

4.2 Derksen and Makam's concativity argument

Here is an algorithmic version of Lemma 2.7 of [DM15].

Lemma 14. *Let $\mathcal{B} \leq M(n, \mathbb{F})$. Assume that for $k, \ell = 1, \ldots, N$ we are given matrices $M_0(k, \ell) \in \mathcal{B}^{\{k,\ell\}}$ of rank $r_0(k, \ell)$, and suppose that $|\mathbb{F}| \geq 2nN + 1$. Then for every $k, \ell = 0, \ldots, N$ we can efficiently (that is, by an algorithm that uses $\mathrm{poly}(Nn)$ arithmetic operations and, over e.g. $\mathbb{Q}$, produces intermediate and final data of size polynomial in the input size) construct matrices $M(k, \ell) \in \mathcal{B}^{\{k,\ell\}}$ of rank $r(k, \ell) \geq r_0(k, \ell)$ such that*

(1) $r(k,\ell+1) \geq r(k,\ell)$ *(*$0 \leq \ell < N$*);*

(2) $r(k+1,\ell) \geq r(k,\ell)$ *(*$0 \leq k < N$*);*

(3) $r(k,\ell+1) \geq \frac{1}{2}(r(k,\ell) + r(k,\ell+2))$ *(*$0 \leq \ell < N-1$*);*

(4) $r(k+1,\ell) \geq \frac{1}{2}(r(k,\ell) + r(k+2,\ell))$ *(*$0 \leq k < N-1$*);*

(5) $r(k,k)$ *is divisible by* k*.*

For $k = 0$ *(resp.* $\ell = 0$*) we assume that* $M_0(k,\ell)$ *is the empty matrix having* ℓ *columns (resp.* k *rows), and* $r(k,\ell) = 0$*.*

Proof. Initially put $M(k,\ell) = M_0(k,\ell)$ for every pair (k,ℓ). For a $k \times \ell$ matrix T let T^+ denote the $(k+1)\times\ell$ matrix obtained form T by appending a zero ($(k+1)$st) row, T^{++} is obtained by appending two zero rows. For $M = \sum_{i=1}^m B_i \otimes T_i$ we use M^+ for $\sum_{i=1}^m B_i \otimes T_i^+$, while $M^{++} = \sum_{i=1}^m B_i \otimes T_i^{++}$.

Let (k,ℓ) be a pair such that any of (1)–(5) is violated. Then we will replace some of the matrices $M(k',\ell')$ with matrices having larger rank. Over an infinite base field like $\mathbb{Q}$, each such replacement step (or each small group consisting of a few them) can be followed by an application of the data reduction procedure from [dGIR96] to keep intermediate (as well as the final) data small.

If (1) is violated then, like in [DM15], replace $M(k+1,\ell)$ with $M(k,\ell)^+$. We can treat a violation of (2) symmetrically.

When (3) is violated we consider the matrix $A = A(t) = M(k+2,\ell) + tM(k,\ell)^{++}$ as a $(k+2)\times\ell$ block matrix consisting of square blocks of size n from $\mathcal{B}$. We can choose t from any subset S of size $2nN+1$ of the base field so that A has rank at least $r(k+2,\ell)$, while the first kn rows form a matrix of rank at least $r(k,\ell)$. This is because a necessary condition for violating either of these two conditions is that the determinant of an appropriate (but unknown) sub-matrix vanishes which determinant is, as a polynomial of degree at most nN in t is not identically zero. The product of these polynomials has degree at most $2nN$ therefore it cannot have more that $2nN$ zeros.

If A has rank larger than $r(k+2,\ell)$ then we replace $M(k+2,\ell)$ with A. Otherwise, like in [DM15], let U be the span of the first kn rows of A, V be the span of the first $(k+1)n$ rows and W be the span of the first kn rows and the last n rows. Note that these collections rows correspond to matrices of the form $A_0 = \sum B_i \otimes T_i$, $A_1 = \sum B_i \otimes T_i'$ and $A_2 = \sum B_i \otimes T_i''$ where T_i are $k \times \ell$ matrices, while T_i' and T_i'' have $(k+1)$ rows and ℓ columns. As $U \leq V \cap W$ and the row space of A is $V + W$, we have $r(k,\ell) \leq \dim U \leq \dim(V \cap W) = \dim V + \dim W - \dim V + W = \dim V + \dim W - r(k+2,\ell)$. It follows that $\dim V + \dim W \geq r(k,\ell) + r(k+2,\ell)$, whence violation of (3) is only possible if either $\dim V$ or $\dim W$ is strictly larger than $\frac{1}{2}(r(k,\ell) + r(k+2,\ell))$. Then we replace $M(k+1,\ell)$ with A_1 or A_2, according to which one has larger rank. A violation of (4) is treated symmetrically.

When (5) is violated then we can apply Lemma 4.

As in each round when violation of (1),. . .,(4) or (5) occurs the rank of at least one of the matrices $M(k,\ell)$ is incremented, the total number of rounds for achieving (1)–(5) is at most $N^3 n$. □

And here is essentially Proposition 2.10 of [DM15]. We include a proof (which is almost literally the same as the proof in [DM15]) here for completeness. We note that this lemma deals only with the property of certain families of functions, without referring to matrices.

Lemma 15 ([DM15, Proposition 2.10]). *Assume that $N > n > 0$, $r : \{0, 1, \ldots, N\}^2 \to \mathbb{Z}$ is a function with $0 \le r(k, \ell) \le \min(k, \ell)n$ for $k, \ell \in \{0, 1, \ldots, N\}$ also satisfying (1)–(5) of Lemma 14. Suppose further that $r(1, 1) > 1$, and there exists d such that $n \le d + 1 \le N$ and $r(d + 1, d + 1) = n(d + 1)$. Then, $r(d, d) = nd$ as well.*

Proof. By $r(d + 1, d + 1) = n(d + 1)$, for $1 \le a < d + 1$,

$$r(d+1, a) \ge \frac{(d+1-1) \cdot r(d+1, 0) + a \cdot r(d+1, d+1)}{d+1} = an.$$

As by assumption $r(d + 1, a) \le an$, we have $r(d + 1, a) = an$. Similarly $r(a, d + 1) = an$ for $1 \le a < d + 1$.

Then we bound $r(1, d)$ as follows:

$$\begin{aligned} r(1,d) &\ge \frac{(d-1) \cdot r(1, d+1) + 1 \cdot r(1,1)}{d} \\ &\ge \frac{(d-1)n + 2}{d} = n - \frac{n-2}{d} > n - 1. \end{aligned}$$

Note that we use $r(1, 1) > 1$ and $d \ge n - 1$. Since $r(1, d) \in \mathbb{Z}$, $r(1, d) = n$.

We are ready to bound $r(d, d)$ then.

$$\begin{aligned} r(d,d) &\ge \frac{(d-1) \cdot r(d+1, d) + 1 \cdot r(1,d)}{d} \\ &= \frac{(d-1)dn + n}{d} = nd - n + \frac{n}{d}. \end{aligned}$$

From $d \ge n - 1$ it is inferred easily that $-n + \frac{n}{d} > -d$. Therefore $nd - n + \frac{n}{d} > (n - 1)d$. By (5) we conclude that $r(d, d) = nd$. □

5 The algorithm

Proof of Theorem 1. Let $B_1, \ldots, B_m$ be the input basis for $\mathcal{B}$. The algorithm is an iteration based on Theorem 5. During the iteration we have a matrix $A = \sum_i B_i \otimes T_i \in \mathcal{B}^{\{d,d\}}$ of rank rd for some integer $d \le r - 1$. We assume that at least one of the basis elements B_i has rank larger than 1, as otherwise [IKS10] works even with $d = 1$. Initially $d = 1$ and A is a basis element having rank at least 2. The procedure behind Theorem 5 either returns an $(n - r)$-shrunk subspace (in which case we are done), or a new matrix (denoted also by A) in a blow-up $\mathcal{B}^{\{d',d'\}}$ of rank at least $(r + 1)d'$ for some $d' < r^2$ together with a square window of size $r + 1$ so that the corresponding sub-matrix of A is of rank $(r + 1)d'$. If $d' > r + 2$ then we apply either of the two methods in Section 4 to this block.

- If Lemma 13 is used, then n in the statement will be $r + 1$, and we shall use it repeatedly to get a matrix in the $(r + 2, r + 2)$-blow-up with a similar content as above.
- If Lemma 14 is used, then n in the statement of the lemma will be $r + 1$, N will be d', $M_0(d', d')$ is the nonsingular $(r+1)d' \times (r+1)d'$ block of A and $M_0(p, q)$ can be actually even the zero matrix for $(p, q) \ne (d', d')$. It will prepare matrices in several not necessarily square blow-ups, among others, most importantly, one in an (r, r)-blow-up. The main content of this are r by r matrices $T_1, \ldots, T_m$ such that the $(r + 1)r \times (r + 1)r$ sub-matrix of $A' = \sum B_i \otimes T_i$ has full rank.

Then we replace A with A' and apply the size reduction procedure of [dGIR96] to arrange that the entries of T_i fall into the prescribed subset of $\mathbb{F}$, and continue the iteration with this new matrix A. □

Acknowledgement. We would like to thank the authors of [GGOW15] and of [DM15] for sharing their ideas with us and making us possible to read early versions of their manuscripts. Part of the work was done when Gábor and Youming were visiting the Centre for Quantum Technologies at the National University of Singapore. Research of the first author was also supported in part by the Hungarian National Research, Development and Innovation Office – NKFIH, Grants NK105645 and 115288. Youming's research was supported by the Australian Research Council DECRA DE150100720.

References

[Der01] Harm Derksen. Polynomial bounds for rings of invariants. *Proceedings of the American Mathematical Society*, 129(4):955–964, 2001.

[dGIR96] 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.

[DM15] Harm Derksen and Visu Makam. Polynomial degree bounds for matrix semi-invariants. preprint ArXiv:1512.03393, 2015.

[FR04] M. Fortin and C. Reutenauer. Commutative/noncommutative rank of linear matrices and subspaces of matrices of low rank. *Séminaire Lotharingien de Combinatoire*, 52:B52f, 2004.

[GGOW15] Ankit Garg, Leonid Gurvits, Rafael Oliveira, and Avi Wigderson. A deterministic polynomial time algorithm for non-commutative rational identity testing. preprint ArXiv:1511.03730, 2015.

[Gur04] Leonid Gurvits. Classical complexity and quantum entanglement. *J. Comput. Syst. Sci.*, 69(3):448–484, 2004.

[IKS10] Gábor Ivanyos, Marek Karpinski, and Nitin Saxena. Deterministic polynomial time algorithms for matrix completion problems. *SIAM J. Comput.*, 39(8):3736–3751, 2010.

[IQS15] Gábor Ivanyos, Youming Qiao, and K. V. Subrahmanyam. Non-commutative Edmonds' problem and matrix semi-invariants. preprint arXiv:1508.00690, 2015.

[Lam91] T.Y. Lam. *A First Course in Noncommutative Rings*. Graduate Texts in Mathematics. Springer, 1991.

[Ram54] K.G. Ramanathan. *Lectures on the Algebraic Theory of Fields*. Tata Institute of Fundamental Research, Bombay, 1954.