Generalized matrix completion and algebraic natural proofs

Markus Bläser*[1], Christian Ikenmeyer†[2], Gorav Jindal‡[2], and Vladimir Lysikov§[1,3]

[1]Department of Computer Science, Saarland University
[2]Max-Planck-Institut für Informatik
[3]Cluster of Excellence MMCI, Saarland University

Abstract. Algebraic natural proofs were recently introduced by Forbes-Shpilka-Volk [FSV17] and independently by Grochow-Kumar-Saks-Saraf [GKSS17] as an attempt to transfer Razborov and Rudichs famous barrier result [RR97] for Boolean circuit complexity to algebraic complexity theory. Razborov and Rudich's barrier result relies on a widely believed assumption, namely, the existence of pseudo-random generators. Unfortunately, there is no known analogous theory of pseudo-randomness in the algebraic setting. Therefore, Forbes et al. use a concept called succinct hitting sets instead. This assumption is related to polynomial identity testing, but it is currently not clear how plausible this assumption is. Forbes et al. are only able to construct succinct hitting sets against rather weak models of arithmetic circuits.

Generalized matrix completion is the following problem: Given a matrix with affine linear forms as entries, find an assignment to the variables in the linear forms such that the rank of the resulting matrix is minimal. We call this rank the completion rank. Computing the completion rank is an NP-hard problem. As our first main result, we prove that it is also NP-hard to determine whether a given matrix can be approximated by matrices of completion rank $\leq b$. The minimum quantity b for which this is possible is called border completion rank (similar to the border rank of tensors). Naturally, algebraic natural proofs can only prove lower bounds for such border complexity measures. Furthermore, these border complexity measures play an important role in the geometric complexity program. As another result we provide a toy setting in which no algebraic natural proofs exist unless $\mathsf{P}^{\#\mathsf{P}} = \exists\mathsf{BPP}$, but where nevertheless geometric complexity theory can prove lower bounds succinctly.

Using our hardness result above, we can prove the following barrier: We construct a small family of matrices with affine linear forms as entries and a bound b, such that at least one of these matrices does not have an algebraic natural proof of polynomial size against all matrices of border completion rank b, unless $\mathsf{coNP} \subseteq \exists\mathsf{BPP}$. This is an algebraic barrier result that is based on a well-established and widely believed conjecture. The complexity class $\exists\mathsf{BPP}$ is known to be a subset of the more well

known complexity class MA in the literature. Thus ∃BPP can be replaced by MA in the statements of all our results. With similar techniques, we can also prove that tensor rank is hard to approximate.

Furthermore, we prove a similar result for the variety of matrices with permanent zero. There are no algebraic polynomial size natural proofs for the variety of matrices with permanent zero, unless $\mathsf{P}^{\#\mathsf{P}} \subseteq \exists\mathsf{BPP}$. On the other hand, we are able to prove that the geometric complexity theory approach initiated by Mulmuley and Sohoni [MS01] yields proofs of polynomial size for this variety, therefore overcoming the natural proofs barrier in this case.

1. Introduction

1.1. Algebraic natural proofs

Algebraic natural proofs were introduced by Forbes, Shpilka, and Volk [FSV17] and independently by Grochow, Kumar, Saks, and Saraf [GKSS17] as an attempt to transfer Razborov and Rudichs famous barrier result [RR97] for Boolean circuit complexity to algebraic complexity theory.

Let X be a set of indeterminates. We fix a set of monomials $\mathcal{M} \subseteq K[X]$ and we consider the linear span $\langle\mathcal{M}\rangle$ of $\mathcal{M}$ in $K[X]$. Every polynomial in $\langle M\rangle$ is of the form $\sum_{m\in\mathcal{M}} c_m m$. Every $f \in \langle M\rangle$ is identified with its list of coefficients $(c_m)_{m\in\mathcal{M}}$. We consider a class $\mathcal{C} \subseteq \langle M\rangle$, we think of $\mathcal{C}$ as the polynomials of "low" complexity in $\langle M\rangle$. An *algebraic proof* or *distinguisher* is a polynomial D in $|\mathcal{M}|$ variables T_m, $m \in M$, that vanishes on the coefficient vectors of all polynomials in $\mathcal{C}$. If for $f \in \langle M\rangle$, $D(f) \neq 0$, then D proves that f is not in $\mathcal{C}$, that is, f has "high" complexity.

Definition 1 (Algebraic Natural Proofs [FSV17, GKSS17])**.** *Let X be a set of variables and let $M \subseteq K[X]$ be a set of monomials. Let $\mathcal{C} \subseteq \langle M\rangle$ be a set of polynomials and let $\mathcal{D} \subseteq K[T_m : m \in M]$.*

A polynomial D is an algebraic $\mathcal{D}$-natural proof against $\mathcal{C}$, *if*

1. *$D \in \mathcal{D}$,*
2. *D is a nonzero polynomial, and*
3. *for all $f \in \mathcal{C}$, $D(f) = 0$, that is, D vanishes on the coefficient vectors of all polynomials in $\mathcal{C}$.*

Furthermore, for $f_0 \in \langle M\rangle$, we call D as above an algebraic $\mathcal{D}$-natural proof for f_0 against $\mathcal{C}$, if we have $D(f_0) \neq 0$. That is, D proves that f_0 is not in $\mathcal{C}$.

A *hitting set* of some class of polynomials $\mathcal{P}$ in μ variables is a set of vectors $\mathcal{H} \subseteq K^\mu$ such that for all $p \in \mathcal{P}$, there is an $h \in \mathcal{H}$ such that $p(h) \neq 0$.

Definition 2 (Succinct hitting sets [FSV17, GKSS17])**.** *Let X be a set of variables and let $M \subseteq K[X]$ be a set of monomials. Let $\mathcal{C} \subseteq \langle M\rangle$ be a set of polynomials and let $\mathcal{D} \subseteq K[T_m : m \in M]$.*

H is a $\mathcal{C}$-succinct hitting set for $\mathcal{D}$ *if*

1. *$H \subseteq \mathcal{C}$ and*

2. H viewed as a set of vectors of coefficients of length $|M|$ is a hitting set for $\mathcal{D}$.

From the definitions above it follows immediately that exactly one of the following is true:

- there is a algebraic $\mathcal{D}$-natural proof against $\mathcal{C}$ or
- $\mathcal{C}$ is a $\mathcal{C}$-succinct hitting set of $\mathcal{D}$.

That is, the existence of succinct hitting sets rules out the existence of natural proofs. Forbes et al. [FSV17] replace succinct hitting set by a more constructive concept called *succinct generators*, but the barrier result stays essentially the same. We refer to their paper for the details.

The most interesting example is when $\mathcal{C}$ is the class of polynomials in n variables that have degree $\mathsf{poly}(n)$ and circuit size $\mathsf{poly}(n)$, that is, we get the class VP when we run over all n (see Section A in the appendix for all relevant background information and [Bür00] for more details). Let N be the number of coefficients of such polynomials. We have $\mathsf{poly}(n) = \mathsf{poly}\log(N)$. An algebraic $\mathsf{poly}(N)$-natural proof is now a polynomial D in N variables that vanishes on $\mathcal{C}$. By the reasoning above, we get the following algebraic natural proofs barrier:

> If there are $\mathsf{poly}\log(N)$ succinct hitting sets for circuits of size $\mathsf{poly}(N)$, then there are no algebraic $\mathsf{poly}(N)$-natural proofs against circuits of size $\mathsf{poly}\log N$.

Forbes et al. construct succinct hitting sets for restricted classes of circuits for which nontrivial lower bounds are known. This might give some evidence, that $\mathsf{poly}\log(N)$ succinct hitting sets for circuits of size $\mathsf{poly}(N)$ might also exists, however, this question is widely open in our opinion.

There is one further problem with the classes studied by Forbes et al.: If a polynomial vanishes on a particular set, it also vanishes on the Zariski closure of this set. So an algebraic proof against some class $\mathcal{C}$ will vanish on polynomials f that are not contained in $\mathcal{C}$, but are contained in the closure $\overline{\mathcal{C}}$. Polynomials in the border $\overline{\mathcal{C}} \setminus \mathcal{C}$ may have higher complexity than polynomials in $\mathcal{C}$ (otherwise, they would be in $\mathcal{C}$), yet they cannot be distinguished by an algebraic proof from polynomials in $\mathcal{C}$, independently of any barrier. Therefore, to study algebraic proofs properly, one needs to look at Zariski closed classes of polynomials. It is important to remark that the complexity of polynomials in the border $\overline{\mathcal{C}} \setminus \mathcal{C}$ may still be polynomially bounded in the complexity of $\mathcal{C}$. Forbes et al. construct succinct hitting sets for many restricted classes of circuits for which nontrivial lower bounds are known. For these circuit classes, it is not known whether they are Zariski closed.

Understanding the border is a fundamental and very difficult problem. In complexity theory it naturally arises in the geometric complexity theory program, see [MS01] and the many subsequent papers as well as [Mul12] for an overview, and the study of tensor rank [BCS97]. Only very little is known about closures and borders. For the exponent of matrix multiplication, see e.g. [Blä13], it does not matter whether one takes rank or border rank as a measure, this is essentially due to the fact that the tensor product of two matrix multiplication tensors is again a matrix multiplication tensor. See also [GMQ16] for some recent progress towards understanding closures.

We refer to the work by Forbes et al. [FSV17] for further details and discussions.

1.2. (Border) tensor rank

Border tensor rank is another application domain for algebraic proofs. It is one of the rare cases where one can show a nontrivial lower bound using the geometric complexity approach [BI13].

One can think of a tensor as a "three-dimensional matrix" $t = (t_{h,i,j}) \in K^{\ell \times m \times n} := K^\ell \otimes K^m \otimes K^n$. A rank-one tensor is a tensor of the form $u \otimes v \otimes w$ with $u \in K^\ell$, $v \in K^m$ and $w \in K^n$. The rank $R(t)$ of t is the smallest number r of rank-one tensors $s_1, \dots, s_r$ such that

$$t = s_1 + s_2 + \cdots + s_r.$$

With each tensor, one can associate a polynomial, for instance a trilinear form

$$t = \sum_{h=}^{\ell} \sum_{i=1}^{m} \sum_{j=1}^{n} t_{h,i,j} X_h Y_i Z_j.$$

So one can view the tensor rank in the above framework of algebraic proofs. However, from the class $\mathcal{C}$, we only use the coefficients (which are the $t_{h,i,j}$). Therefore, we can work with the tensors directly. Let

$$S_r = \{s \in K^\ell \otimes K^m \otimes K^n \mid R(s) \leq r\}$$

be the set of all tensors of rank at most r. An algebraic proof that $R(t) > r$ is a polynomial P in ℓmn variables such that P vanishes on S_r and $P(t) \neq 0$. However, the set S_r is not Zariski-closed. That is, it is not the vanishing set of a set of polynomials. So we look at the Zariski closure X_r of S_r instead. These tensors are called the tensors of border rank $\leq r$. Tensors that are sitting in the border of X_r, that is, in $X_r \setminus S_r$, have some rank $r' > r$. However, there will be no algebraic proof for $R(t') \geq r'$, since any polynomial P that vanishes on $S_{r'-1}$ vanishes also on X_r, since $S_r \subseteq S_{r'-1}$ and hence, $X_r = \overline{S_r} \subseteq \overline{S_{r'-1}}$. Therefore $P(t) = 0$. So the appropriate quantity to study when considering algebraic proofs is the border rank.

Forbes et al. [FSV17] discuss (border) tensor rank briefly at the end of Section 1.2. We will define a related quantity, the so-called (border) completion rank. For the border completion rank, we will prove that there are no algebraic natural proofs of polynomial size, unless $\mathsf{coNP} \subseteq \exists\mathsf{BPP}$.

Given a tensor t and a bound b, it is NP-hard to decide whether $R(T) \leq b$ as shown by Håstad [Hås90], see also the work by Shitov [Shi16] and Schaefer and Stefankovic [SS16] for improvements. It is not known whether the same is true for the border rank. However, we can show that is true for border completion rank.

1.3. Matrix completion problems

An instance of a matrix completion problem over some field K is an $n \times n$-matrix A that is filled with elements from K or with a special symbol $*$. One can think of the $*$'s as placeholders that can be replaced by arbitrary elements from K. The goal is to replace the $*$'s in such a way that the rank of the resulting matrix is either minimized or maximized, depending on the application.

Matrix completion has many applications, for instance, in machine learning and network coding, we here just refer to [Pee96,HKY06,HMRW14], which contain relevant hardness results. When we consider minimization, the problem is NP-hard, even when the resulting matrix has rank 3 [Pee96]. When we consider maximization, then the problem is NP-hard over finite fields [HKY06]. Over large enough fields, there is a simple randomized polynomial time algorithm that simply works by plugging in random elements from a large enough set. The correctness of this algorithm follows from the well-known Schwartz-Zippel lemma. Derandomising this matrix completion algorithm is a major open problem, however, quite recently a deterministic quasi-polynomial time (even quasi-NC) algorithm was given by Gurjar and Thierauf [GT17].

We can phrase the matrix completion problem as a problem on tensors. Let $E_{i,j} \in K^{n\times n}$ be the matrix that has a 1 in position (i,j) and zeros elsewhere. Let A_0 be the matrix that is obtained from A by replacing every $*$ by a 0. For every star, we create a matrix $E_{i,j}$ where (i,j) is the position of the $*$. Let $F_1,\dots,F_m$ be the resulting matrices. We can view $(A_0,F_1,\dots,F_m)$ as a tensor in $K^{n\times n\times(m+1)}$. We call $A_0,F_1,\dots,F_m$ the slices of this tensor. Then the matrix completion problem can be phrased as follows: Find the minimum r such that there are $\lambda_1,\dots,\lambda_m \in K$ fulfilling

$$\mathrm{rk}(A_0+\lambda_1F_1+\cdots+\lambda_mF_m)\le r.$$

Here rk denotes the usual matrix rank. Many variants of matrix completion have been studied in the literature. For instance, instead of having simply $*$'s we can have variables instead and each occurrence of a variable has to be replaced by the same value. This can naturally be modeled as a tensor problem, too: Each of the F_i will have a 1 at each position where a particular variable occurs and 0's elsewhere. The most general setting would be the following: Given a tensor t as a tuple of $n\times n$-matrices $(A_0,A_1,\dots,A_m)$, what is the minimum r such that there are $\lambda_1,\dots,\lambda_m$ with

$$\mathrm{rk}(A_0+\lambda_1A_1+\cdots+\lambda_mA_m)\le r.$$

We call this problem a generalized matrix completion problem and we call the minimum value r above the *completion rank* of t.

We can view an instance of a matrix completion problem also as a matrix with affine linear forms in variables $x_1,\dots,x_m$ as entries. We will freely switch between these two representation, a tensor with slices $(A_0,A_1,\dots,A_m)$ or a matrix $A_0+x_1A_1+\dots x_mA_m$ with affine linear forms in variables $x_1,\dots,x_m$ as entries.

1.4. Our contribution

We study tensors $t=(A_0,A_1,\dots,A_m)$ given by $m+1$ slices of size $n\times n$. This is a tensor with $n^2(m+1)$ many entries. We are interested in the class of all tensors with completion rank bounded by some given r. We prove that given a tensor t and a bound r, deciding whether the completion rank of t is bounded by r is NP-hard. This might not be astonishing, since the same is true for the ordinary tensor rank. Then, we define the border completion rank: t has border completion rank $\le r$ if t is contained in the Zariski closure of the set of all tensors of completion rank $\le r$ (where the closure is taken in some appropriately chosen variety, see Section 2 for more details). We go on by showing that it is even NP-hard to check given t and r, whether the border completion rank of t is bounded by r, that is, whether t is contained in the closure of the set of all tensors with completion rank $\le r$. Completion rank is therefore one of the rare examples where we understand the border. Formally, we prove the following theorem in Section 3.

Theorem 3. *Let K be a field of characteristic distinct from 2. Given a tensor t and an integer r, deciding whether the border completion rank $\underline{\mathrm{CR}}(t)\le r$ is NP-hard.*

Hardt, Meka, Raghavendra, and Weitz [HMRW14] prove the hardness of some kind of approximate matrix completion problem. Our result is fundamentally different. In their setting, the size of all values are bounded. For defining border complexity, one needs to consider unbounded entries!

Next we construct a small family of tensors (small means that they come from a closed, even low dimensional set) such that not all of these tensors can have algebraic $\mathsf{poly}(n)$-natural[1] proofs against the set of all tensors of completion rank $\leq r$ for some appropriately chosen r. This means that there is a tensor t such that any polynomial D with $D(t) \neq 0$ that vanishes on all tensor of completion rank r has super-polynomial circuit complexity. This results is of course conditional, but it is based on the widely believed condition that $\mathsf{coNP} \not\subseteq \exists\mathsf{BPP}$. More specifically, we prove the following theorem in Section 3.

Theorem 4. *For infinitely many n, there is an m, a tensor $t \in K^{n\times n\times m}$ with coefficients in $\{-1, 0, 1\}$, and a value r such that there is no algebraic $\mathsf{poly}(n)$-natural proof for the fact that $\underline{\mathrm{CR}}(t) > r$ unless $\mathsf{coNP} \subseteq \exists\mathsf{BPP}$.*

One can view this as a meta-result: Proving lower bounds via algebraic proofs is difficult. At least, if we want to represent the proof by an algebraic circuit. Note that even the geometric complexity approach eventually produces an algebraic natural proof. However, it is produced from some intermediate representation, which might be more compact. Barrier results for the geometric complexity program seem even harder to obtain.

We can also use our constructions to prove that it is NP-hard to approximate the tensor rank up to a factor $(1+\epsilon)$ for some $\epsilon > 0$. This means in particular, that there is no FPTAS for tensor rank, unless $\mathsf{P} = \mathsf{NP}$, see Section 5. It was pointed out to us by one of the STOC 2018 referees that this result was proven independently by Song, Woodruff, and Zhong [SWZ17] by modifying Håstad's construction. Our construction is simpler.

Grochow and Pitassi [GP14] study so-called ideal proof systems. They use their framework, among many other things, to give a short proof of a transfer theorem, namely that $\mathsf{VP}^0 = \mathsf{VNP}^0$ implies that $\mathsf{coNP} \subseteq \exists\mathsf{BPP}$. We could do a similar proof using our barrier result. However, we notice that one gets an equally short proof for an even stronger transfer theorem by considering polynomials that vanish on the set of matrices with permanent equal to 0. This result in formally proved in Section 6 as stated below.

Theorem 5. *If $\mathsf{VP}^0 = \mathsf{VNP}^0$, then $\mathsf{P}^{\#\mathsf{P}} \subseteq \exists\mathsf{BPP}$.*

Using basically the same proof we relate this to algebraic natural proofs as follows.

Theorem 6. *If there are VP^0-natural proofs against the set of matrices with permanent zero, then $\mathsf{P}^{\#\mathsf{P}} \subseteq \exists\mathsf{BPP}$.*

In Section 7 we show how so-called *occurrence obstructions* in geometric complexity theory can be used to decide whether the permanent of a matrix is zero, thus breaking the natural proofs barrier.

Theorem 7. *There is a polynomial sized occurrence obstruction for the set of matrices with nonzero permanent.*

The reader is referred to Section 7 for a short introduction to geometric complexity theory and the necessary definitions.

[1] Note that the tensors we consider have $n^2(m+1)$ many entries. But we can bound $m \leq n^2$, since otherwise, there will be a linear dependence between $A_1, \dots, A_m$ and we can shrink t without changing its (border) completion rank.

1.5. Further related work

Grochow and Pitassi [GP14] introduce algebraic proof systems, which are different from algebraic natural proofs. They study so-called ideal proof systems, which are based on Hilbert's Nullstellensatz. Such a proof is a polynomial C that proves that polynomials $F_1, \dots, F_m$ do not have a common zero. The inputs to C are not the coefficients of $F_1, \dots, F_m$ but $F_1, \dots, F_m$ themselves.

Efremenko, Garg, Oliveira, and Wigderson [EGdOW17] study limits of rank-based methods (so-called flattenings) for proving lower bounds on tensor rank and Waring rank (of tensors in $K^{\otimes d}$).

2. Generalized matrix completion problems

We shall use K to denote the underlying field over which we consider the matrix completion instances. One can think of $K = \mathbb{C}$.

Definition 8. *Let $A_0, A_1, \dots, A_m \in K^{n\times n}$. The* completion rank *of $A_0, A_1, \dots, A_m$ is the minimum number r such that there are scalars $\lambda_1, \dots, \lambda_m$ with*

$$\mathrm{rk}(A_0 + \lambda_1 A_1 + \cdots + \lambda_m A_m) \le r.$$

We denote the completion rank by $\mathrm{CR}(A_0, A_1, \dots A_m)$.

The set of all $(m+1)$-tuples of $n \times n$-matrices together with m scalars $\lambda_1, \dots, \lambda_m$

$$(A_0, A_1, \dots, A_m, \lambda_1, \dots, \lambda_m) \in K^{(m+1)n^2+m}$$

such that

$$\mathrm{rk}(A_0 + \lambda_1 A_1 + \dots \lambda_m A_m) \le r$$

is a closed set, since it is defined by vanishing of all $(r+1)\times(r+1)$-minors. Denote this set by $P_r^{m,n}$. (We will omit the m and n if they are clear from the context.) We can also view $(A_0, A_1, \dots, A_m)$ as a tensor in $K^{n\times n\times(m+1)}$ with slices $A_0, A_1, \dots, A_m$.

Let $C_r^{m,n}$ be the projection of $P_r^{m,n}$ onto the first $(m+1)n^2$ components, that is, $C_r^{m,n}$ is the set of all $(A_0, A_1, \dots, A_m)$ with $\mathrm{CR}(A_0, A_1, \dots, A_m) \le r$. Note that $C_r^{m,n}$ need not be closed. Indeed, consider

$$A_0 = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} \qquad \text{and} \qquad A_1 = \begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix}.$$

Clearly, $\mathrm{CR}(A_0, A_1) = 2$. But we have

$$\underbrace{\begin{pmatrix} 1 & 0 \\ \epsilon & 1 \end{pmatrix}}_{=:A_{0,\epsilon}} + \frac{1}{\epsilon}\begin{pmatrix} 0 & 1 \\ 0 & 0 \end{pmatrix} = \begin{pmatrix} 1 & 1/\epsilon \\ \epsilon & 1 \end{pmatrix}.$$

(Think of ϵ being a small number for the moment.) Thus $\mathrm{CR}(A_{0,\epsilon}, A_1) = 1$ for every $\epsilon \neq 0$. Then $(A_{0,\epsilon}, A_1)$ converges to (A_0, A_1) in the Euclidean topology. Hence (A_0, A_1) has completion rank 2 but is contained in the Euclidean (and also Zariski) closure of C_1.

In our definition of border completion rank, an important question is: with respect to which field shall we take the Zariski closure? Let B be any rank-one matrix. Then the completion rank of (I, B) is at least $n-1$. Here, I is the $n \times n$ identity matrix. We can approximate B by $B + \epsilon I$. But $I - \frac{1}{\epsilon}(B + \epsilon I)$ has rank 1 and in fact, this trick always works. Therefore, it seems reasonable that the rank of the approximating matrices should be the same as the matrix itself, so we take the closure in $K^{n\times n} \times K^{n\times n}_{r_1} \times \cdots \times K^{n\times n}_{r_m}$, where $K^{n\times n}_{\rho}$ denotes the closed set of matrices of rank at most ρ and $r_i = \mathrm{rk}(A_i)$.

In the following, we give an algebraic definition of border completion rank, as was done for the tensor rank, see [BCS97]. This has the advantage that it is independent of the underlying field. One can also give a definition in terms of limits; over $\mathbb{C}$, these two notions coincide. In the following, let ϵ denote some indeterminate. We will now consider our tensors over the field $K(\epsilon)$ of rational functions. For $f, g \in K(\epsilon)$, we write $f = g + O(\epsilon^i)$ if the coefficients of f and g agree for powers ϵ^j with $j < i$ when expanded as formal Laurent series (around 0). We write $A = B + O(\epsilon^i)$ for two matrices A and B with entries from $K(\epsilon)$, when the same is true in every component.

Definition 9. *Let $A_0, A_1, \ldots, A_m \in K^{n\times n}$. The* border completion rank *of $A_0, A_1, \ldots, A_m$ is the minimum number r such that there are approximations $\tilde{A}_i \in K(\epsilon)^{n\times n}$ with $\tilde{A}_i = A_i + O(\epsilon)$, $0 \le i \le m$ and rational functions $\lambda_1, \ldots, \lambda_m \in K(\epsilon)$ with*

$$\mathrm{rk}(\tilde{A}_0 + \lambda_1 \tilde{A}_1 + \cdots + \lambda_m \tilde{A}_m) \le r.$$

We denote the border completion rank by $\underline{\mathrm{CR}}(A_0, A_1, \ldots A_m)$.

See Section B for a discussion of alternative ways to define the closure and border completion ranks.

3. Border completion rank and natural proofs

Let ϕ be a formula in 2-CNF over the variables $x_1, \ldots, x_t$ with clauses $c_1, \ldots, c_s$. We want to use NP-hardness of the Max-2-SAT problem to prove that both completion rank and border completion rank are NP-hard. More specifically, the following problem Max-2-SAT is NP-hard, see [ACG+99]: Given a formula ϕ in 2-CNF and a bound $b \in \mathbb{Z}^+$, decide whether there is an assignment to the variables of ϕ that satisfies at least b clauses of ϕ.

We will define an instance of border matrix completion of size $n = 2s$ and $m = t$. Our matrices will have a block structure, there will be s blocks of size 2×2, one for each clause.

In the actual construction, the clause gadgets will be on the diagonal of some larger block diagonal matrix. The constants will appear in the 0^{th} layer of the tensor that we construct and the coefficients of the i^{th} variable will appear in the i^{th} layer.

Let $c_i = L_1 \vee L_2$ be a clause in ϕ. The corresponding clause gadget looks like

$$\begin{pmatrix} 1-\ell_1 & 1 \\ 0 & 1-\ell_2 \end{pmatrix}$$

Here ℓ_j in the matrix is x_k if the literal $L_j = x_k$ and it is $1 - x_k$ if $L_j = \neg x_k$, $j = 1, 2$. All these clause gadgets are blocks of our desired block diagonal matrix. More specifically, take these s clause gadgets as above, one for each clause, and form a block diagonal matrix of it. We get a matrix with affine linear forms as entries. Write this matrix as $A_0 + x_1 A_1 + \ldots x_t A_t$. $(A_0, A_1, \ldots, A_t)$ is our matrix completion instance.

Observation 10. *The clause gadget has rank* 1 *iff at least one of the literals ℓ_1, ℓ_2 is set to be* 1*. Otherwise, it has rank* 2*.*

By using above observations, following lemma follows immediately.

Lemma 11. *$\mathrm{CR}(A_0, A_1, \ldots, A_t) \leq 2s - b$ iff b clauses of ϕ can be satisfied. Thus the problem $\mathrm{CR}(A_0, A_1, \ldots, A_t) \overset{?}{\leq} k$ is* NP*-hard.*

Now our goal is to prove Lemma 11 for $\underline{\mathrm{CR}}$ also. Since all variables appear only on the diagonals of the gadgets and only one variable in each entry, we observe:

Observation 12. *If $i \geq 1$, then each A_i is a diagonal matrix with diagonal entries being ± 1. Moreover, if the j^{th} diagonal entry of A_i is non-zero then the j^{th} diagonal entry of any other A_k is zero, for $i, k \geq 1$.*

Let $\tilde{A}_0, \tilde{A}_1, \ldots, \tilde{A}_t$ be approximations to $A_0, A_1, \ldots, A_t$, that is, $\tilde{A}_i = A_i + O(\epsilon)$.

Lemma 13. *There are (invertible) matrices $S = I_n + O(\epsilon)$ and $T = I_n + O(\epsilon)$ such that $S \cdot (\tilde{A}_0 + \lambda_1 \tilde{A}_1 + \cdots + \lambda_t \tilde{A}_t) \cdot T = \hat{A}_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t$ for some $\hat{A}_0 = A_0 + O(\epsilon)$.*

Proof. We first show the existence of S_1, T_1 of the form $I_n + O(\epsilon)$ such that $S_1 \tilde{A}_1 T_1 = A_1$. Let $j_1, \ldots, j_p$ be the indices such that the positions $(j_1, j_1), \ldots, (j_p, j_p)$ contain exactly the ± 1's of A_1. Using elementary row and column operations, we can achieve that $\tilde{A}_1$ agrees with A_1 in all columns and rows $j_1, \ldots, j_p$. Note that these elementary and column operations are of following two forms.

1. Dividing a row from rows $j_1, \ldots, j_p$ by some constant of the form $1 + O(\epsilon)$.

2. Adding an $O(\epsilon)$ multiple of some k^{th} row or column to some other row or column, where $k \in \{j_1, \ldots, j_p\}$.

Each of the above operations corresponds to left or right multiplying by a matrix of the form $I_n + O(\epsilon)$. Thus there exist matrices S_1 and T_1 of the form $I_n + O(\epsilon)$, such that $S_1 \tilde{A}_1 T_1 = B_1$, where B_1 agrees with A_1 in all columns and rows $j_1, \ldots, j_p$. Since S_1, T_1 are of full rank, we get $\mathrm{rk}\, A_1 = \mathrm{rk}\, \tilde{A}_1$, thus B_1 cannot have any nonzero entries outside these rows $j_1, \ldots, j_p$ and columns $j_1, \ldots, j_p$. Therefore, $S_1 \tilde{A}_1 T_1 = B_1 = A_1$.

Similarly, we obtain S_2, T_2 of the form $I_n + O(\epsilon)$ such that $S \tilde{A}_2 T = A_2$. By observing the form of elementary and column operations in S_2 and T_2 as above and also using the observation above that non-zero diagonal entries of A_1 and A_2 occur at distinct indices, we get that $S_2 A_1 T_2 = A_1$. Now we can continue this process of converting $\tilde{A}_i$ to A_i as well. Thus we may assume that the approximation $\tilde{A}_i$ are replaced by the exact matrices A_i. This process changes $\tilde{A}_0$ to some other approximation of A_0, which we called $\hat{A}_0$ in the statement of the lemma. □

Since S and T above have full rank, we get that $\mathrm{rk}(\tilde{A}_0 + \lambda_1 \tilde{A}_1 + \cdots + \lambda_t \tilde{A}_t) = \mathrm{rk}\,(\hat{A}_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t)$. Thus to prove the NP-hardness of border completion rank, we can assume that $\tilde{A}_i = A_i$ for $i \geq 1$. We rename $\hat{A}_0$ back to $\tilde{A}_0$. Assume there are $\lambda_i = a_{i,0} \epsilon^{d_i} + a_{i,1} \epsilon^{d_i + 1} + \ldots$ with $a_{i,0} \neq 0$ such that

$$\mathrm{rk}(\tilde{A}_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t) \leq 2s - b. \tag{1}$$

Here, d_i can be any integer, not necessarily non-negative.

Lemma 14. *$\underline{\mathrm{CR}}(A_0, A_1, \ldots, A_t) \leq 2s - b$ iff b clauses of ϕ can be satisfied.*

Proof. By Lemma 11, we know that if b clauses of ϕ can be satisfied then $\mathrm{CR}(A_0, A_1, \ldots, A_t) \leq t + 2s - b$, thus $\underline{\mathrm{CR}}(A_0, A_1, \ldots, A_t) \leq \mathrm{CR}(A_0, A_1, \ldots, A_t) \leq 2s - b$.

Now we show that if $\underline{\mathrm{CR}}(A_0, A_1, \ldots, A_t) \leq 2s - b$ then b clauses of ϕ can be satisfied. This means that there exist λ_i as in (1).

If all $d_i \geq 0$, then we can substitute $\epsilon = 0$ in (1) to get $\mathrm{rk}(A_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t) \leq 2s - b$, thus $\mathrm{CR}(A_0, A_1, \ldots, A_t) \leq 2s - b$. By again using Lemma 11, we obtain that b clauses of ϕ can be satisfied. Thus we can assume w.l.o.g. that some $d_k < 0$. Note that the λ_i induce an assignment to the x_i and thus to literals ℓ_j's in the clauses, although this assignment might not be Boolean and might have negative powers of ϵ. A clause gadget of clause $c_i = L_1 \vee L_2$ in $\tilde{A}_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t$ looks like

$$\begin{pmatrix} 1 + O(\epsilon) - \ell_1 & 1 + O(\epsilon) \\ O(\epsilon) & 1 + O(\epsilon) - \ell_2 \end{pmatrix}$$

For above clause gadget to have rank one, we need to have $\ell_1 = 1 + O(\epsilon)$ or $\ell_2 = 1 + O(\epsilon)$. We call such clauses to be "ϵ-satisfied" which satisfy $\ell_1 = 1 + O(\epsilon)$ or $\ell_2 = 1 + O(\epsilon)$. Suppose we have at least b "ϵ-satisfied" clauses. By substituting $\epsilon = 0$ in corresponding λ_i's, we get an assignment of corresponding x_i's. The rest of the x_i's are assigned arbitrary 0 or 1 values. It is clear that this assignment of x_i's satisfies all the clauses which were "ϵ-satisfied" , thus existence of b "ϵ-satisfied" clauses implies that b clauses of ϕ can be satisfied.

Therefore we can assume that there are less than b "ϵ-satisfied" clauses. We shall use this assumption to prove that $\mathrm{rk}(\tilde{A}_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t) > 2s - b$. For this, we shall construct a non-vanishing minor M of size at least $2s - b + 1$.

We take both rows and columns of clause gadgets for which the corresponding clause is not "ϵ-satisfied". For the clause gadgets whose clause is "ϵ-satisfied", we take the first row and second column, this is entry $1 + O(\epsilon)$. This completes the construction of M. Since there are less than b "ϵ-satisfied" clauses, size of the M is $> 2s - b$. Let us use k to denote the size of M, so we have $k > 2s - b$.

Now we show that M does not vanish. If $M = (m_{i,j})$ then $\det(M) = \sum_{\sigma \in S_k} \mathrm{sign}(\sigma) \prod_{i \in [k]} m_{i,\sigma(i)}$. By construction, the elements on the diagonal of M have the unique power of ϵ of minimum degree in their respective row. Thus this power of ϵ can not be canceled by any other product term in the sum $\sum_{\sigma \in S_k} \mathrm{sign}(\sigma) \prod_{i \in [k]} m_{i,\sigma(i)}$. Therefore $\det(M) \neq 0$. Hence $\mathrm{rk}(\tilde{A}_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t) > 2s - b$, a contradiction. □

Proof of Theorem 3. This follows immediately from Lemma 14 and the NP-hardness of Max-2-SAT. □

Let $t \in K^{m \times n \times (m+1)}$ be a tensor. Recall that an *algebraic* $\mathsf{poly}(n)$*-natural proof* for the border completion rank of t being $> r$ is a polynomial equation $p \in K[X_{h,i,j} | 1 \leq i, j \leq n,\ 1 \leq j \leq m]$ such that

1. $p(t) \neq 0$,
2. $p(s) = 0$ for every $s \in K^{n \times n \times (m+1)}$ with $\underline{\mathrm{CR}}(s) \leq r$.
3. p is computed by a constant-free algebraic circuit of size $\mathsf{poly}(n)$.

We here confine ourselves to constant-free circuits, we can deal with arbitrary constants similar to [GP14]. Note that our proofs have an additional condition (compared to Forbes et al. [FSV17]) that it shall vanish on a certain tensor t. So we want a concrete proof for some tensor t. (Note however, that if $p(t) \neq 0$, then p proves this for a lot of tensors, since the set at which p does not vanish is open.) Forbes et al. only demand that p is nonzero, that is, that it can prove a lower bound for unknown some tensor (and henceforth for many unknown tensors).

Observation 15. *Let $U_{i,j}, V_{i,j}$, $1 \leq i \leq \rho$, $1 \leq j \leq n$ be indeterminates. Consider the polynomial matrix $\sum_{i=1}^{\rho}(U_{i,1}, \dots, U_{i,n})^T(V_{i,1}, \dots, V_{i,n})$. If we substitute arbitrary constants for the indeterminates, then we get all matrices in $K_\rho^{n\times n}$*

Lemma 16. *Let $Q_0, Q_1, \dots, Q_t$ be polynomial matrices as in Observation 15, having ranks $r_0, \dots, r_t$, respectively. We use fresh variables for each Q_i. Let $g := (Q_0 - Z_0Q_1 - \cdots - Z_tQ_t, Q_1, \dots, Q_t)$, where $Z_1, \dots, Z_t$ are new variables. If we substitute arbitrary constants for the indeterminates, then we get all tensors of completion rank $\leq r_0$ with the i^{th} slice having rank $\leq r_i$, $1 \leq i \leq t$.*

Proof. By Observation 15, the images of each Q_i give all of the matrices of rank r_i. If a tensor $(A_0, A_1, \dots, A_t)$ has completion rank at most r_0, then then there are scalars $\lambda_1, \dots, \lambda_t$ such that $A_0 + \lambda_1 A_1 + \cdots + \lambda_t A_t$ has rank $\leq r_0$. This means that there is a rank r_0 matrix B such that $A_0 = B - \lambda_1 A_1 - \cdots - \lambda_t A_t$. This is exactly generated in the first component of g. □

Proof of Theorem 4. The tensors will be the tensors constructed above from 2-CNF formulas. Assume that the assertion of the theorem is false and let n be large enough such that for all tensors under consideration, there is a algebraic $\mathsf{poly}(n)$-natural proof.

Let ϕ be a formula in 2-CNF and let $b \in \mathbb{N}$. We want to check whether every assignment satisfies $< b$ clauses of ϕ. This problem is coNP-hard. Let $T_\phi = (A_0, \dots, A_t)$ be the corresponding tensor constructed above in Theorem 3. Let p be a polynomial that vanishes on all tensors of border completion rank $\leq 2s - b$ but not on T_ϕ and that has polynomial-size arithmetic circuits. Such a polynomial is guaranteed to exist by assumption. However, we do not know how to construct such a p.

We can use the nondeterminism to guess a polynomial-sized circuit. Then we have to verify that the polynomial p computed by the circuit has indeed this property.

Let g be the tensor of Lemma 16 with $r_0 = 2s - b$ and $r_i = \mathrm{rk}(A_i)$, $1 \leq i \leq t$. If $p(g) = 0$ as a polynomial identity, then by Lemma 16, p vanishes on all tensors of completion rank r_0 that have ranks of the slices $1, \dots, t$ bounded by $r_1, \dots, r_t$. Since p is a polynomial, it also vanishes on all tensors of border completion rank r_0 that have ranks of the slices $1, \dots, t$ bounded by $r_1, \dots, r_t$.

Now given T_ϕ, we can decide whether b clauses of ϕ cannot be satisfied, as follows:

1. Guess a circuit C of polynomial size computing a polynomial p.
2. Decide whether $p(g) = 0$ using polynomial identity testing.
3. Check whether $p(T_\phi) \neq 0$. If yes, then accept. Otherwise reject.

The correctness follows from the construction. It is obviously an $\exists\mathsf{BPP}$ algorithm. Note that $p(T_\phi) \neq 0$ can again be checked by polynomial identity testing. (A direct evaluation might not be possible, since this could involve large numbers.) □

Note that the nonexistence of algebraic $\mathsf{poly}(n)$-natural proofs does not simply follow from the NP-hardness of border completion rank. It could have been the case that all the proofs have polynomial size, but they cannot be constructed in polynomial time. Or there could be many such proofs, because the variety has many components.

Observation 17. *In the hardness proof for border completion rank, we look at tensors* $(A_0, A_1, \dots, A_t)$ *of border completion rank at most* r *with* $\mathrm{rk}(A_i) \leq c$ *for all* i *for some constant* c *and each* A_i *has size* $n := 2s$. *The variety* $C_r^{n,t,c}$ *of all such tensors is irreducible: We can think of* g *constructed in Lemma 16 as a polynomial map from* $K^{n\times r} \times K^{r\times n} \times (K^{n\times c} \times K^{c\times n})^t$ *to* $K^{n\times n\times t}$. *The closure of the image of* g *is* $C_r^{n,t,c}$. *Therefore,* $C_r^{n,t,c}$ *itself is irreducible.*

We can strengthen the statement of Theorem 4: There are infinite sequences $t_n = \Theta(n)$ and $r_n = \Theta(n)$ and a constant c such that for every set of equations describing the variety $C_{r_n}^{n,t_n,c}$, at least one equation has superpolynomial circuit complexity, unless $\mathsf{coNP} \subseteq \exists\mathsf{BPP}$.

4. Relation to (border) tensor rank

The following theorem was essentially shown by Derksen [Der14]. The proof uses the well-known fact that if a slice of a tensor t has rank-one, then we can use this slice in an optimal decomposition of t into rank-one tensors.

Theorem 18. *If* $t = (A_0, A_1, \dots, A_m)$ *is tensor such that* $A_0, A_1, \dots, A_t$ *are linearly independent and* $\mathrm{rk}(A_1) = \cdots = \mathrm{rk}(A_m) = 1$. *Then*

$$\mathrm{R}(t) = \mathrm{CR}(t) + m.$$

It is not clear whether the same is true for border rank and border completion rank. We can prove that the right-hand side is an upper bound for the left-hand side.

Proposition 19. *If* $t = (A_0, A_1, \dots, A_m)$ *is a tensor such that* $\mathrm{rk}(A_1) = \cdots = \mathrm{rk}(A_m) = 1$. *Then*

$$\underline{\mathrm{R}}(t) \leq \underline{\mathrm{CR}}(t) + m.$$

Proof. Let $r = \underline{\mathrm{CR}}(t)$ and let $(\tilde{A}_0, \tilde{A}_1, \dots, \tilde{A}_m)$ be approximations to $(A_0, A_1, \dots, A_t)$ with $\mathrm{rk}(\tilde{A}_i) = \mathrm{rk}(A_i) = 1$ and $\lambda_1, \dots, \lambda_m \in K(\epsilon)$ such that

$$\mathrm{rk}(\tilde{A}_0 + \lambda_1 \tilde{A}_1 + \cdots + \lambda_m \tilde{A}_m) \leq r.$$

We can write each $\tilde{A}_i = u_i \otimes v_i$ with $u_i, v_i \in K(\epsilon)^n$. Since $\mathrm{rk}(\tilde{A}_0 + \lambda_1 \tilde{A}_1 + \cdots + \lambda_m \tilde{A}_m) \leq r$, there are vectors $x_1, \dots, x_r, y_1, \dots, y_r \in K(\epsilon)$ such that

$$\tilde{A}_0 + \lambda_1 \tilde{A}_1 + \cdots + \lambda_m \tilde{A}_m = x_1 \otimes y_1 + \cdots + x_r \otimes y_r.$$

Therefore,

$$t = e_0 \otimes x_1 \otimes y_1 + \cdots + e_0 \otimes x_r \otimes y_r + (-\lambda_1 e_0 + e_1) \otimes u_1 \otimes v_1 + \cdots + (-\lambda_m e_0 + e_m) \otimes u_m \otimes v_m + O(\epsilon),$$

where $e_0, \dots, e_m$ denotes the standard basis (corresponding to the slices). Therefore, $\underline{\mathrm{R}}(t) \leq r + m$. ∎

Can we prove an analogue of Theorem 4 for the border rank? First of all, we do not know whether a converse of Theorem 18 is true. More severe, however, seems to be the restriction to matrices of rank 1. In the next Section 5, we will introduce variable and clause gadgets that only contain rank-one matrices. But they only work for the tensor rank, not for the border tensor rank, since they allow solutions in the border that do not correspond to Boolean assignments. Border completion rank is more forgiving; there we could use more compact gadgets, that did not allow unwanted solutions in the border, because there were no rank restrictions.

5. Tensor rank is hard to approximate

There is a $\delta > 0$ and a constant c such that the following promise problem is NP-hard: Given a formula ϕ in 3-CNF such that every variable appears in at most c clauses with the promise that either ϕ is satisfiable or any assignment can satisfy at most a fraction of $1 - \delta$ of the clauses, decide which is the case, see e.g. [ACG+99, Theorem 8.13]. Using this as the starting point for the proof of Theorem 3, we get the following result:

Theorem 20. *Let K be any field. There is a constant $\gamma > 0$ such that given a tensor t, it is* NP*-hard to approximate the completion rank of t within a factor of $(1 + \gamma)$.*

Proof. Let ϕ be a formula in 3-CNF with t variables and s clauses such that each variable appears in at most c clauses. We construct a matrix complextion instance $(A_0, A_1, \dots, A_t)$ similarly to the corresponding (border) completion rank instance in the proof of Theorem 3. A_0 has s blocks of size 3×3 for the clauses, since we now have a 3-SAT instance. The clause gadget for a clause $c = L_1 \vee L_2 \vee L_3$ looks like

$$\begin{pmatrix} 1 - \ell_1 & 1 & 0 \\ 0 & 1 - \ell_2 & 1 \\ 0 & 0 & 1 - \ell_3 \end{pmatrix}.$$

where $\ell_1 = x_i$ if x_i appears positively in the first literal L_1 of the clause and $1 - x_i$ otherwise. ℓ_2 and ℓ_3 are defined accordingly. The gadget has the same properties as the previous gadget for clauses of length 2: If we set the variables such that at least one ℓ_i becomes one, then the gadget has rank two. For any other assignment, it has rank three. In particular, setting a variable to any other value than zero or one can never reduce the rank of the gadget. Therefore, we do not need any variable gadgets.

If ϕ is satisfiable, then $\mathrm{CR}(A_0, A_1, \dots, A_t) = 2s$; this rank is achieved by any satisfying assignment to ϕ.

Now assume that ϕ is not satisfiable, then any assignment to the variables can satisfy at most $(1 - \delta)s$ clauses. Since we are only considering completion rank, the situation is much easier compared to border completion rank: $A_0 + \lambda_1 A_1 + \dots + \lambda_t A_t$ is a block diagonal matrix. The rank of $A_0 + \lambda_1 A_1 + \dots + \lambda_t A_t$ is $2s + u$ where u is the number of clause gadgets that have rank three. From $\lambda_1, \dots, \lambda_t$, we construct a Boolean assignment by setting x_i to λ_i if $\lambda_i \in \{0, 1\}$ and to an arbitrary Boolean value otherwise. This Boolean assignment satisfies at least $s - u$ clauses. Since ϕ is not satisfiable, $u \geq \delta s$. Therefore, $\mathrm{CR}(A_0, A_1, \dots, A_t) \geq (2 + \delta)s$.

Therefore, the reduction is approximation preserving and we obtain the statement of the theorem. □

Next, we want to use this construction to show that also tensor rank is hard to approximate. The idea is to go via Theorem 18. To this aim, we have to ensure that every matrix $A_1, \dots, A_t$ has rank 1. Our construction is inspired by the work of Schaefer and Stefankovic [SS16], but it is tailored to formulas in 3-CNFs and therefore it is easier to analyze, in particular when proving approximation hardness.

Assume we have a clause with variables x, y and z. The clause gadget looks as follows:

$$\begin{pmatrix} 1 & x & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 1 & u & 0 & 0 & 0 & 0 & s(u)-u_1 & 0 & 0 \\ 0 & 0 & 1 & y & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & v & 0 & 0 & 0 & s(v)-v_1 & 0 \\ 0 & 0 & 0 & 0 & 1 & z & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & w & 0 & 0 & s(w)-w_1 \\ 0 & u-u_2 & 0 & 0 & 0 & 0 & 1-\ell(u) & 1 & 0 \\ 0 & 0 & 0 & v-v_2 & 0 & 0 & 0 & 1-\ell(v) & 1 \\ 0 & 0 & 0 & 0 & 0 & w-w_2 & 0 & 0 & 1-\ell(w) \end{pmatrix}.$$

Next comes the clause gadget. We have

$$\ell(u) = \begin{cases} u & \text{if } x \text{ appears positive in the clause,} \\ 1-u & \text{otherwise,} \end{cases}$$

and

$$s(u) = \begin{cases} -u & \text{if } x \text{ appears positive in the clause,} \\ u & \text{otherwise.} \end{cases}$$

$\ell(v)$, $s(v)$, $\ell(w)$, and $s(w)$ are defined accordingly.

All matrices have rank one. The gadget consists of three 2×2-blocks on the diagonal which create local variables u, v, and z, which are copies of x, y, and z. The 3×3-block in the lower right corner implements the clause. All other entries not in these blocks are only introduced to make the matrices rank 1.

The variables u, v, and w are local variables that are only used within the gadget. The variables u_i, v_i, and w_i are also local variables, they are needed to make all occurring matrices rank one: For instance, the local variable u appears on the diagonal twice, locally in positions $(2,2)$ and $(7,7)$. To make the corresponding 2×2-submatrix rank one, we add the variable u to the two other corners $(2,7)$ and $(7,2)$ of this 2×2-submatrix. The entry position $(7,7)$ is either u or $-u$. In order to get a rank one matrix, the entry in position $(2,7)$ will also be u or $-u$, respectively. To be able to recover the original diagonal submatrix, we also add new local variables u_1 and u_2 in the positions $(2,7)$ and $(7,2)$. The same is done for the variables v and w.

The variables x, y and z (and all further variables of the formula, of course), appear in several clause gadgets. Therefore, the matrix that corresponds to say x does not have rank one. We use the same construction as for the local variables in the overall construction, which is described after the next lemma.

Lemma 21. *1. If we set x, y, z to values from $\{0,1\}$ such that the clause is satisfied, then the local variables in the clause gadget can be set such that the resulting matrix has rank five.*

2. *If the variables are set in such a way that the rank of the clause gadget is five, then x, y, z are set such that the clause is satisfied, which means that at least one variable is set to a value from $\{0, 1\}$ and this value satisfies the corresponding literal.*

Proof. For the first item, we construct an assignment for the local variables as follows: we set $u = x$, $v = y$, and $w = z$ and all variables u_i, v_i, and w_i in such a way that the unique entry they appear in becomes zero. In this way, we are left with a block diagonal matrix. In the three 2×2 blocks, the second row equals the first row or is zero. The 3×3-block is the original clause gadget. Since the assignment satisfies the clause, the rank of this block is two. Therefore, the overall rank is five.

For the second item, we will derive several equations from the fact that the rank of the gadget is five. Consider the clause gadget with rows 4,6,9 and columns 4,6,7 removed:

$$\begin{pmatrix} 1 & x & 0 & 0 & 0 & 0 \\ 1 & u & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & u - u_2 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 - \ell(v) & 1 \end{pmatrix}.$$

Its determinant is $u - x$. Next consider the gadget with rows 2,6,9 and columns 2,6,7 removed:

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & y & 0 & 0 & 0 \\ 0 & 1 & v & 0 & s(v) - v_1 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & v - v_2 & 0 & 1 - \ell(v) & 1 \end{pmatrix}.$$

Its determinant is $v - y$. Next comes the gadget with rows 2,4,9 and columns 2,4,7 removed:

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & z & 0 & 0 \\ 0 & 0 & 1 & w & 0 & s(w) - w_1 \\ 0 & 0 & 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 - \ell(v) & 1 \end{pmatrix}.$$

The determinant is $w - z$. Finally consider the gadget with rows and columns 2,4,6 removed:

$$\begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 - \ell(u) & 1 & 0 \\ 0 & 0 & 0 & 0 & 1 - \ell(v) & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 - \ell(w) \end{pmatrix}.$$

Its determinant is $(1 - \ell(u))(1 - \ell(v))(1 - \ell(w))$.

Assume that we have an assignment such that the clause gadget has rank 5. Then the vanishing of the first three minors enforces that $x = u$, $y = v$, and $z = w$ and the vanishing of the last one that at least one literal is one and hence the clause is satisfied (and therefore at least one of x, y or z is from $\{0, 1\}$). □

Let ϕ be a formula in 3-CNF with t variables and s clauses such that every variable appears in a constant number c of clauses. Note that $s = O(t)$. In our final construction, for every clause, we will have one clause gadget. These gadgets are arranged as a block diagonal matrix. The matrices that correspond to an actual variables of ϕ are not rank-one matrices. We make them rank-one as we did locally in the clause gadgets, as described above: If x appears in positions $(i_1, j_1), \dots, (i_k, j_k)$, then we replace this matrix by a rank one matrix with 1's in all positions (i_h, j_ℓ), $1 \leq h, \ell \leq k$. k is at most c, since each variable appears in at most c clauses. For all tuples (i_h, j_ℓ) with $h \neq \ell$, we have an additional rank-one matrix with a 1 in this position and 0s elsewhere, corresponding to a new variable $x_{h,\ell}$ in the completion problem. In this way, the original diagonal matrix is in the span of these matrices. Let T_ϕ be the resulting tensor. By construction, all slices have rank 1 except for the 0th slice. The number of slices of T_ϕ is $O(t)$, since each clause gadget introduces a constant number of slices. Furthermore, there are only a constant number, namely $\leq c^2$, of additional rank-one matrices introduced in the last step.

Lemma 22. *Assume that ϕ is either satisfiable or any assignment satisfies at most $(1 - \epsilon)$ of the clauses for some $\epsilon > 0$.*

1. *If ϕ is satisfiable, then the completion rank of T_ϕ is at most $5s$.*

2. *If ϕ is not satisfiable, then the completion rank of T_ϕ is at least $5s + \epsilon s$.*

Proof. If ϕ is satisfiable, then we set each variable in T_ϕ corresponding to an original variable of ϕ to the value of a satisfying assignment. By Lemma 21, we can set all clause gadgets locally such that their rank is 5. We set the auxiliary variables $x_{h,\ell}$ in such a way that we clear all elements not in a block on the diagonal. Therefore, the completion rank is at most $5s$.

For the second item, assume that the completion rank of T_ϕ is $r < 5s + \epsilon s$. Choose an assignment to the variables of T_ϕ that achieves this completion rank. Choose linear independent columns $a_1, \dots, a_r$ of the resulting matrix A.

From every clause gadget, the local columns 1, 3, 5, 8, and 9 are linearly independent, even if we remove the local rows 1, 3, and 5. In the overall construction, there are only nonzero entries outside the gadget in the local columns 2, 4, and 6 and rows 1, 3, and 5. Therefore, we can always assume that for any clause gadget, its local columns 1, 3, 5, 8, and 9 are among the $a_1, \dots, a_r$.

Now consider a gadget G that contributes more than five columns to $a_1, \dots, a_r$. There are at most ϵs such gadgets. Assume that the local column 2 of G is among $a_1, \dots, a_r$ and it is a_i. Furthermore, let the variable in a_i be x. Then we can use the variables $x_{h,j}$ to clear all entries in the column a_i outside the gadget. This will not affect the other columns of A, since every column has its own variables. If thereafter, a_i is in the span of the remaining columns, then we can exchange it against a new column from A, since r was the completion rank of T_ϕ. By repeating this process if necessary, we can assume that no a_i has nonzero entries outside the clause gadgets, since the same reasoning works for the local columns 4 and 6 and the local column 7 has no nonzero entries outside the gadget by construction. Therefore, we can assume w.l.o.g. that A is a block diagonal matrix and the blocks on the diagonal are instantiations of the clause gadgets.

Now we have clause gadgets with five columns among the $a_1, \dots, a_r$ and clause gadgets with six or more. We claim that the assignment to the variables of T_ϕ that correspond to variables of ϕ is an assignment that satisfies all clauses with only five columns among $a_1, \dots, a_r$. If this is the case, then we get a satisfying assignment which satisfies a fraction of more than $(1-\epsilon)$ of all clauses and henceforth, ϕ is satisfiable.

Let G be a clause gadget with only five columns among $a_1, \dots, a_r$. Choose $r-5$ rows from the other clause gadgets to get an invertible $(r-5)\times(r-5)$-submatrix M of A, which is disjoint from G. Choose a 6×6-submatrix S of G. As a submatrix of A, S and M form a block diagonal matrix. The determinant of this matrix vanishes, since its size is $(r+1)\times(r+1)$, but r is the completion rank. Since M is invertible, it follows that all 6×6 minors of G vanish. By Lemma 21, the assignment satisfies the corresponding clause. □

Theorem 23. *Tensor rank is* NP*-hard to approximate.*

Proof. Let T_ϕ be the tensor in Lemma 22. By Theorem 18, $R(T_\phi) = \mathrm{CR}(T_\phi) + k$ where k is the total number of slices of T_ϕ. Since every variable of ϕ occurs only a constant number of times in ϕ and every gadget has constant size, $k = \Theta(t)$. Since $\mathrm{CR}(T_\phi)$ is hard to approximate, so is $R(T_\phi)$. □

6. Matrices with permanent zero and a transfer theorem

In this section we prove Theorem 5 and 6. When Grochow and Pitassi [GP14] studied so-called ideal proof systems, they used their framework to provide a short proof of a transfer theorem, namely that $\mathsf{VP}^0 = \mathsf{VNP}^0$ implies that $\mathsf{coNP} \subseteq \exists\mathsf{BPP}$. Such a transfer theorem shows that separations in the Boolean world yield separations in the algebraic world. (We do not know of any transfer theorem in the other direction.) One could use our framework to do something similar; however, with roughly the same effort, we can prove an even stronger transfer theorem. The starting point is algebraic VP^0-natural proofs for the set of all matrices A, say with rational entries, that fulfill $\mathrm{Per}(A) = 0$. (Note that we allow negative entries.) Note that this set has a VNP^0-natural proof, namely the permanent itself. We will prove that if there is a family of polynomials $(D_n) \in \mathsf{VP}^0$ such that D_n is nonzero and vanishes on all matrices with permanent 0, then we can construct a circuit for the permanent in $\exists\mathsf{BPP}$. This uses the self-reducibility of the permanent, similar to [KI04], and Kaltofen's factorization algorithm [Kal89]. In particular, if $\mathsf{VP}^0 = \mathsf{VNP}^0$, then this will allow us to show that $\mathsf{P}^{\#\mathsf{P}} \subseteq \exists\mathsf{BPP}$.

Let X be an $n\times n$ matrix. Construct an $n\times n$ matrix Z as follows:

$$\begin{cases} z_{ij} = x_{ij} & \text{for } i \le n-1, \\ z_{nj} = x_{nj}\,\mathrm{Per}\,X_{nn} & \text{for } j \le n-1, \\ z_{nn} = -\sum_{j=1}^{n-1} x_{nj}\,\mathrm{Per}\,X_{nj}, \end{cases}$$

where X_{ij} is the matrix obtained from X by removing the i^{th} row and the j^{th} column. We have $\mathrm{Per}\,Z = 0$. Moreover, any matrix with $\mathrm{Per}\,Z = 0$ and $\mathrm{Per}\,Z_{nn} \neq 0$ can be obtained in this way, by Laplace expansion. The following theorem from [Kal89] states that if a polynomial can be computed by small circuit then so can be its factors. We only need this result for prime fields, so we state it only for prime fields $\mathbb{F}_p$.

Theorem 24 (Kaltofen [Kal89])**.** *If $f \in \mathbb{F}_p[x_1, \dots, x_n]$ can be computed by an arithmetic circuit of size at most s and degree at most d, $f = \prod_{i=1}^{k} g_i^{p^{e_i}\cdot j_i}$, where the g_i's are irreducible and $p \nmid j_i$ for each i. Then we can compute, for each $i \in [k]$, the numbers e_i, j_i, and an arithmetic circuit for the factor $g_i^{p^{e_i}}$ in randomized* $\mathsf{poly}(n, s, d, \log p)$ *time.*

Although theorem [Kal89] works for the factors of the form $g_i^{p^{e_i}}$, in our special case e_i would always be zero, thus giving us the factor g_i directly. We also need the following lemma from [FPdV13], which bounds the size of coefficients of any VP^0 polynomial.

Lemma 25 (Fournier-Perifel-Verclos [FPdV13])**.** *Let P be a polynomial computed by an arithmetic circuit of size s and formal degree d with constants of absolute value bounded by $M \geq 2$, then the sum of the absolute values of its coefficients is at most $M^{s \cdot d}$.*

Now we are ready to prove our result.

Proof of Theorem 5. If $\mathsf{VP}^0 = \mathsf{VNP}^0$, there exists a VP^0-family of circuits computing Per_n. Let $s(n)$ be a polynomial bound on the size of a circuit for Per_n and $d(n)$, on its formal degree.

We describe an $\exists\mathsf{BPP}$ algorithm which computes a circuit for Per_n over $\mathbb{F}_p[x_{1,1}, \ldots, x_{n,n}]$ for some $p > n!$. These circuits can be used to implement any language from $\mathsf{P}^{\#\mathsf{P}}$.

The algorithm uses non-determinism to guess a circuit for Per_k for each $k \leq n$. We start from the trivial circuit for Per_1. For each k, perform the following steps:

1. Use non-determinism to guess a VP^0 circuit C_k of size $\leq s(k)$, which potentially computes Per_k.

2. Check that its formal degree is $\leq d(k)$, otherwise reject.

3. Use a randomized algorithm for PIT to check that $C_k(Z_k(X)) = 0$, otherwise reject. Here $Z_k(X)$ denotes the matrix constructed above of size $k \times k$ with entries taken from $X = (x_{i,j})$. Note that the parameterization Z_k uses Per_{k-1}, which can be computed by the circuit for Per_{k-1} obtained on the previous iteration. Here we make sure that this PIT succeeds with probability at least $1 - \frac{1}{3n^2}$.

4. By Hilbert's Nullstellensatz, the polynomial computed by C_k has the form $(\mathrm{Per}_k)^e h$ with $e \leq \frac{d(k)}{k}$ and $\gcd((\mathrm{Per}_k)^e, h) = 1$. Note that Per_k is irreducible.

5. Since C_k has a VP^0 circuit of size $s(k)$ and formal degree is $d(k)$, we know that the magnitude of coefficients of C_k is bounded by $B := 2^{d(k) \cdot s(k)}$, this follows from Lemma 25. So we choose a prime $p > \max\{B, n!\}$. By the choice of p, we see that $C_k = (\mathrm{Per}_k)^e h$ even over $\mathbb{F}_p[x_{1,1}, \ldots x_{k,k}]$.

6. Now we use Theorem 24 to obtain a circuit for all the irreducible factors g_i of $C_k = \prod_{i=1}^{k} g_i^{p^{e_i} \cdot j_i}$ in randomized $\mathsf{poly}(n, s(k), d(k), \log p) = \mathsf{poly}(n, s(k), d(k))$ time. Since we choose p to be large enough, we know that all $e_i = 0$ and $p \nmid j_i$ for all i. Here also, we assume that this randomized algorithm succeeds with probability at least $1 - \frac{1}{3n^2}$.

7. For all irreducible factors g_i's of C_k, use a randomized algorithm for PIT to check that $g_i(Z_k(X)) = 0$. Here we make sure that this PIT succeeds with probability at least $1 - \frac{1}{3 \cdot d(k) \cdot n^2}$. If this PIT succeeds then we know that g_i has to be Per_k since g_i is irreducible. Thus we can compute a constant free circuit for Per_k with probability $1 - \frac{1}{n^2}$, of size $\mathsf{poly}(n, s(k), d(k))$.[2]

[2]Note that we factor a newly guessed circuit every time. We only need the circuit of the previous round to check whether C_k vanishes on $Z_k(X)$ but not as a building block of C_k. So we will not have a repeated increase in circuit size due to factoring.

Since we have n iterations, overall success probability is at least $1 - \frac{1}{n}$. Thus $\mathsf{P}^{\#\mathsf{P}} \subseteq \exists\mathsf{BPP}$. $\square$

A slight modification of the proof of Theorem 5 gives a proof of Theorem 6 as follows.

Proof of Theorem 6. Let the set of permanent 0 matrices have VP^0-natural proofs. Then there exists a family $(C_k)_{k\in\mathbb{N}} \in \mathsf{VP}^0$ of nonzero polynomials such that C_k vanishes on $k \times k$ matrices of permanent 0. Let $s(k)$ be the size of C_k and $d(k)$ be its formal degree. Both are polynomially bounded by definition of VP^0.

We proceed precisely as in the proof of Theorem 5, but in step 1 we guess a circuit C_k of size $\leq s(k)$ with the property that it computes a nonzero polynomial that vanishes on $k \times k$ matrices of permanent 0. The non-deterministic guess C_k is not required to compute the permanent correctly. $\square$

7. Geometric complexity theory breaks the natural proofs barrier

In this section we show how geometric complexity theory breaks the natural proofs barrier presented in Theorem 6.

Geometric complexity theory is an approach towards algebraic complexity lower bounds via tools from algebraic geometry and representation theory. In [MS01, MS08] Mulmuley and Sohoni define *obstructions*, which are representation theoretic multiplicities that can be used to show complexity lower bounds. In this section we discuss how this approach can potentially break the algebraic natural proofs barrier. Indeed, we will show that it breaks the barrier in the "permanent = 0" example.

Let $Z \subseteq \mathbb{C}^{n\times n}$ be the variety of matrices with permanent zero. Theorem 6 shows that if Z has VP^0-natural proofs, then $\mathsf{P}^{\#\mathsf{P}} \subseteq \exists\mathsf{BPP}$. In this section we show that for every point outside of Z we can prove that it lies outside of Z by using representation theoretic *occurrence obstructions*, analogously to those that were proposed in [MS01, MS08], see the definitions below. Since these occurrence obstructions have a succinct encoding, in the "permanent = 0" setting geometric complexity theory breaks the algebraic natural proof barrier.

For the necessary background on group actions and representation theory of the symmetric group, the general linear group, and the algebraic torus we direct the reader to the lecture notes [BI17], or to the excellent textbooks [Sag01], [Ful97, parts I and II], and [FH91, Ch. 4 and Ch. 15]. The group $\mathrm{GL}_n \times \mathrm{GL}_n$ acts on the set of matrices $\mathbb{C}^{n\times n}$ via left-right multiplication:

$$(g_1, g_2) \cdot A := g_1 A (g_2)^T,$$

for $A \in \mathbb{C}^{n\times n}$, $(g_1, g_2) \in \mathrm{GL}_n \times \mathrm{GL}_n$, where the transpose is taken for technical reasons[3]. Let $T_n \subseteq \mathrm{GL}_n$ denote the algebraic torus (i.e., the group of diagonal matrices with nonzero determinant) and let $\mathfrak{S}_n \subseteq \mathrm{GL}_n$ denote the symmetric group embedded into GL_n via permutation matrices. Let $Q_n \subseteq \mathrm{GL}_n$ denote the group of monomial matrices, i.e., matrices with nonzero determinant that have a single nonzero entry in each row and column. As a group Q_n is isomorphic to a semi-direct product $T_n \rtimes \mathfrak{S}_n$. Since the permanent is invariant (up to scale) under rescaling of rows and columns and permuting rows and columns, the variety Z is closed under the action of the group $G := Q_n \times Q_n \subseteq \mathrm{GL}_n \times \mathrm{GL}_n$, which means that if $A \in Z$, then $gA \in Z$ for all $g \in G$.

[3] The transpose gives $g(g'A) = (gg')A$ for all pairs $g \in \mathrm{GL}_n \times \mathrm{GL}_n$, $g' \in \mathrm{GL}_n \times \mathrm{GL}_n$.

To show that a matrix A does not lie in Z, the geometric complexity approach goes as follows. For the sake of contradiction, assume that $A \in Z$. Then the whole orbit $GA := \{gA \mid g \in G\}$ is contained in Z, because Z is closed under the action of G. Since Z is also Zariski-closed, we have $\overline{GA} \subseteq Z$ as a sub-variety. Let $\mathbb{C}[\mathbb{C}^{n\times n}] := \mathbb{C}[x_{1,1}, \ldots, x_{n,n}]$ denote the $\mathbb{C}$-algebra of polynomials on the space of $n \times n$-matrices. For a Zariski-closed subset $Y \subseteq \mathbb{C}^{n\times n}$ let $I(Y) \subseteq \mathbb{C}[\mathbb{C}^{n\times n}]$ denote the *vanishing ideal* of Y, i.e., the ideal of polynomials that vanish identically on Y. If Y is invariant under rescaling (which is true for our sets of interest, $Y = Z$ or $Y = \overline{GA}$), then the vanishing ideal inherits the grading from the polynomial ring: We denote by $I(Y)_d$ the homogeneous degree d component of $I(Y)$. We define the *coordinate ring* $\mathbb{C}[Y]$ of Y as the quotient $\mathbb{C}[Y] := \mathbb{C}[\mathbb{C}^{n\times n}]/I(Y)$, which also naturally inherits the grading $\mathbb{C}[Y]_d := \mathbb{C}[\mathbb{C}^{n\times n}]_d/I(Y)_d$. Since $\overline{GA} \subseteq Z$, it follows $I(Z)_d \subseteq I(\overline{GA})_d$ for all d. Therefore the restriction of functions gives a canonical surjection between finite dimensional vector spaces:

$$\mathbb{C}[Z]_d \twoheadrightarrow \mathbb{C}[\overline{GA}]_d. \tag{2}$$

Since Z and $\overline{GA}$ are closed under the action of G, we get an action on the coordinate rings via the so-called *canonical pullback*

$$(gf)(B) := f(g^T B)$$

for $g \in G$, $f \in \mathbb{C}[Y]$, $B \in \mathbb{C}^{n\times n}$. In this way, both $\mathbb{C}[Z]_d$ and $\mathbb{C}[\overline{GA}]_d$ are G-representations in the following sense:

For a group H (e.g. $H = G$, $H = T_n$, $H = \mathrm{GL}_n$, $H = T_n \times T_n$, $H = \mathfrak{S}_n \times \mathfrak{S}_n$), an *$H$-representation* is a finite dimensional vector space V with a group homomorphism $\varrho : H \to \mathrm{GL}(V)$. For an element $g \in H$ and a vector $f \in V$ we use the shorthand notation gf for $(\varrho(g))(f)$. A linear map $\varphi : V_1 \to V_2$ between two H-representations V_1 and V_2 is called *equivariant* if for all $g \in H$ and $f \in V_1$ we have $\varphi(gf) = g\varphi(f)$. A bijective equivariant map is called an H-isomorphism. Two H-representations are called *isomorphic* if an H-isomorphism exists from one to the other. A linear subspace of an H-representation that is closed under the action of H is called a *subrepresentation*. An H-representation whose only subrepresentations are itself and 0 is called *irreducible*.

For many groups H the irreducible representations have a complete classification. In order to state Proposition 26, we give now describe a natural index set for all irreducible G-representations. A *partition* λ of n is a non-increasing list of positive integers that sum up to n. To each irreducible representation of G we can assign a *type*, which is a pair of tuples in which each tuple consists of a list of n integers and a partition of n. Two irreducible G-representations are isomorphic iff they have the same type.

Let (1^n) denote the list of n many 1's. We fix the type $\nu := (((1^n),(n)),((1^n),(n)))$. The irreducible G-representation V of type ν is 1-dimensional, $V = \langle f \rangle$, where the symmetric groups act trivially and the tori act by rescaling:

$$\big(\pi, \mathrm{diag}(\alpha_1, \ldots, \alpha_n)\ ;\ \sigma, \mathrm{diag}(\beta_1, \ldots, \beta_n)\big) f = \alpha_1 \cdots \alpha_n \beta_1 \cdots \beta_n f.$$

We will use the classification of irreducible representations for other groups H later in the proof of Proposition 26.

The group G (as well as the other groups that we consider in the proof of Proposition 26) is *linearly reductive*, which means that every G-representation V decomposes into a direct sum of irreducible representations. This decomposition is not necessarily unique, but for each type λ the *multiplicity* $\mathrm{mult}_\lambda(V)$ of λ in V is unique, which is the number of summands that have type λ. The famous Schur's lemma says for an equivariant map $\varphi : V \to W$, the image $\varphi(V)$ is a G-representation and multiplicities do not increase when applying φ:

$$\text{mult}_\lambda(V) \geq \text{mult}_\lambda(\varphi(V)).$$

Recall $A \in Z$. The map in (2) is equivariant and thus

$$\text{mult}_\lambda(\mathbb{C}[Z]_d) \geq \text{mult}_\lambda(\mathbb{C}[\overline{GA}]_d). \tag{3}$$

A type λ that violates (3) is called an *obstruction*. By the above reasoning, the existence of an obstruction proves $A \notin Z$. In other words,

$$\text{if there exists } \lambda \text{ with } \text{mult}_\lambda(\mathbb{C}[Z]_d) < \text{mult}_\lambda(\mathbb{C}[\overline{GA}]_d), \text{ then } A \notin Z. \tag{4}$$

Proposition 26. *Let $G := Q_n \times Q_n$. We fix the type $\nu := (((1^n),(n)),((1^n),(n)))$.*
We have

- $\text{mult}_\nu(\mathbb{C}[Z]_n) = 0$ *and*
- $\text{mult}_\nu(\mathbb{C}[\overline{GA}]_n) = \begin{cases} 0 & \text{if } A \in Z, \\ 1 & \text{otherwise.} \end{cases}$

Proposition 26 says that if $A \notin Z$, then this can be shown using (4), which can be thought of as a succinct presentation of the permanent. Here we are in the special case where one of the multiplicities is zero. These obstructions are called *occurrence obstructions* in the literature.

Proof. The irreducible representations of GL_n are indexed by partitions of length at most n and denoted by $\{\lambda\}$. The irreducible representations of $\mathfrak{S}_n$ are indexed by partitions of n and denoted by $[\lambda]$. The group $\text{GL}_n \times \text{GL}_n \times \mathfrak{S}_d \times \mathfrak{S}_d$ is linearly reductive and its irreducibles are tensor products $\{\lambda\} \otimes \{\mu\} \otimes [\nu] \otimes [\xi]$. The famous Schur-Weyl duality states that the dth tensor power decomposes as follows as a $\text{GL}_n \times \text{GL}_n \times \mathfrak{S}_d \times \mathfrak{S}_d$-representation:

$$\bigotimes^d(\mathbb{C}^n) \otimes \bigotimes^d(\mathbb{C}^n) = \bigoplus_{\lambda,\mu \vdash_n d} \{\lambda\} \otimes \{\mu\} \otimes [\lambda] \otimes [\mu],$$

where $\lambda, \mu \vdash_n d$ means that we sum over pairs of partitions of d that each have length at most n. We have the natural isomorphism

$$\bigotimes^d(\mathbb{C}^n) \otimes \bigotimes^d(\mathbb{C}^n) \simeq \bigotimes^d(\mathbb{C}^{n\times n}).$$

We symmetrize with respect to $\mathfrak{S}_d$ to convert tensors into polynomials:

$$\begin{aligned} \mathbb{C}[\mathbb{C}^{n\times n}]_d &= \bigoplus_{\lambda,\mu \vdash_n d} \{\lambda\} \otimes \{\mu\} \otimes ([\lambda] \otimes [\mu])^{\mathfrak{S}_d} \\ &= \bigoplus_{\lambda \vdash_n d} \{\lambda\} \otimes \{\lambda\}, \end{aligned}$$

where $([\lambda] \otimes [\mu])^{\mathfrak{S}_d}$ denotes the $\mathfrak{S}_d$-invariant space in $[\lambda] \otimes [\mu]$, which is 1-dimensional iff $\lambda = \mu$, and 0-dimensional otherwise. The irreducible representations of T_n are indexed by lists of n integers. Now choose $d = n$ and consider the irreducible $T_n \times T_n$ subrepresentation of type $((1^n),(1^n))$, where we use Gay's theorem [Gay76] that the sum of irreducible T_n-representations of type (1^n) in $\{\lambda\}$ is isomorphic to $[\lambda]$ as an $\mathfrak{S}_n$-representation, provided λ is a partition of n: $\{\lambda\}_{(1^n)} = [\lambda]$:

$$(\mathbb{C}[\mathbb{C}^{n\times n}]_n)_{((1^n),(1^n))} \quad = \quad \bigoplus_{\lambda \vdash n} [\lambda] \otimes [\lambda].$$

As an $\mathfrak{S}_n \times \mathfrak{S}_n$-representation, the trivial subrepresentation (i.e., the representation of type $((n),(n))$) occurs once in this decomposition. This subrepresentation is 1-dimensional, so as a G-representation we have

$$\begin{aligned} W := (\mathbb{C}[\mathbb{C}^{n\times n}]_n)_\nu &= \mathbb{C}, \\ \mathrm{mult}_\nu(\mathbb{C}[\mathbb{C}^{n\times n}]_n) &= 1. \end{aligned}$$

But the permanent polynomial lies in W, so W is the line spanned by the permanent. Since the permanent vanishes on Z we have $\mathrm{mult}_\nu(I(Z)_n) = 1$ and thus $\mathrm{mult}_\nu(\mathbb{C}[Z]_n) = 0$. For $A \in Z$ we have $\overline{GA} \subseteq Z$, therefore $\mathrm{mult}_\nu(\mathbb{C}[\overline{GA}]_n) = 0$. For every matrix A that does not lie in Z we have that the permanent does not vanish on A, so $\mathrm{mult}_\nu(I(\overline{GA}))_n) = 0$ and therefore $\mathrm{mult}_\nu(\mathbb{C}[\overline{GA})]_n = 1)$. □

Now Theorem 7 immediately follows from Proposition 26. Note that the type ν in the proposition is just an efficient encoding of the permanent. So while it is a short proof, its verification is hard. But this is to be expected, since the permanent is a hard function.

8. Acknowledgement

We would like to thank the STOC 2018 referees for their detailed and very helpful comments. We would also like to thank Peter Bürgisser for enlightening discussions at the Workshop on Algebraic Complexity Theory (WACT 2018) in Paris.

References

[ACG+99] G. Ausiello, P. Crescenzi, G. Gambosi, V. Kann, A. Marchetti-Spaccamela, and M. Protasi. *Complexity and Approximation.* Springer, 1999.

[BCS97] Peter Bürgisser, Michael Clausen, and Mohammad Amin Shokrollahi. *Algebraic complexity theory*, volume 315 of *Grundlehren der mathematischen Wissenschaften.* Springer, 1997.

[BI13] Peter Bürgisser and Christian Ikenmeyer. Explicit lower bounds via geometric complexity theory. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *Symposium on Theory of Computing Conference, STOC'13, Palo Alto, CA, USA, June 1-4, 2013*, pages 141–150. ACM, 2013.

[BI17] Markus Bläser and Christian Ikenmeyer. Introduction to geometric complexity theory. lecture notes: `https://people.mpi-inf.mpg.de/~cikenmey/teaching/summer17/introtogct/gct.pdf`, 2017.

[Blä13] Markus Bläser. Fast matrix multiplication. *Theory of Computing, Graduate Surveys*, 5:1–60, 2013.

[Bür00] P. Bürgisser. *Completeness and Reduction in Algebraic Complexity Theory*. Algorithms and Computation in Mathematics. Springer Berlin Heidelberg, 2000.

[Der14] Harm Derksen. On the equivalence between low rank matrix completion and tensor rank. *CoRR*, abs/1406.0080, 2014.

[EGdOW17] Klim Efremenko, Ankit Garg, Rafael Mendes de Oliveira, and Avi Wigderson. Barriers for rank methods in arithmetic complexity. *CoRR*, abs/1710.09502, 2017.

[FH91] William Fulton and Joe Harris. *Representation Theory - A First Course*, volume 129 of *Graduate Texts in Mathematics*. Springer, 1991.

[FPdV13] Hervé Fournier, Sylvain Perifel, and Rémi de Verclos. *On Fixed-Polynomial Size Circuit Lower Bounds for Uniform Polynomials in the Sense of Valiant*, pages 433–444. Springer Berlin Heidelberg, Berlin, Heidelberg, 2013.

[FSV17] Michael A. Forbes, Amir Shpilka, and Ben Lee Volk. Succinct hitting sets and barriers to proving algebraic circuits lower bounds. In Hatami et al. [HMK17], pages 653–664.

[Ful97] William Fulton. *Young tableaux*, volume 35 of *London Mathematical Society Student Texts*. Cambridge University Press, Cambridge, 1997.

[Gay76] D.A. Gay. Characters of the Weyl group of $SU(n)$ on zero weight spaces and centralizers of permutation representations. *Rocky Mountain J. Math.*, 6(3):449–455, 1976.

[GKSS17] Joshua A. Grochow, Mrinal Kumar, Michael E. Saks, and Shubhangi Saraf. Towards an algebraic natural proofs barrier via polynomial identity testing. *CoRR*, abs/1701.01717, 2017.

[GMQ16] Joshua A. Grochow, Ketan D. Mulmuley, and Youming Qiao. Boundaries of VP and VNP. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 34:1–34:14. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2016.

[GP14] Joshua A. Grochow and Toniann Pitassi. Circuit complexity, proof complexity, and polynomial identity testing. In *55th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2014, Philadelphia, PA, USA, October 18-21, 2014*, pages 110–119. IEEE Computer Society, 2014.

[GT17] Rohit Gurjar and Thomas Thierauf. Linear matroid intersection is in quasi-NC. In Hatami et al. [HMK17], pages 821–830.

[Hås90] Johan Håstad. Tensor rank is NP-complete. *J. Algorithms*, 11(4):644–654, 1990.

[HKY06] Nicholas J. A. Harvey, David R. Karger, and Sergey Yekhanin. The complexity of matrix completion. In *Proceedings of the Seventeenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2006, Miami, Florida, USA, January 22-26, 2006*, pages 1103–1111. ACM Press, 2006.

[HMK17] Hamed Hatami, Pierre McKenzie, and Valerie King, editors. *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*. ACM, 2017.

[HMRW14] Moritz Hardt, Raghu Meka, Prasad Raghavendra, and Benjamin Weitz. Computational limits for matrix completion. In Maria-Florina Balcan, Vitaly Feldman, and Csaba Szepesvári, editors, *Proceedings of The 27th Conference on Learning Theory, COLT 2014, Barcelona, Spain, June 13-15, 2014*, volume 35 of *JMLR Workshop and Conference Proceedings*, pages 703–725. JMLR.org, 2014.

[Kal89] Erich Kaltofen. Factorization of polynomials given by straight-line programs. In *Randomness and Computation*, pages 375–412. JAI Press, 1989.

[KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing polynomial identity tests means proving circuit lower bounds. *Comput. Complex.*, 13(1/2):1–46, December 2004.

[MS01] Ketan Mulmuley and Milind A. Sohoni. Geometric complexity theory I: an approach to the P vs. NP and related problems. *SIAM J. Comput.*, 31(2):496–526, 2001.

[MS08] K.D. Mulmuley and M. Sohoni. Geometric Complexity Theory. II. Towards explicit obstructions for embeddings among class varieties. *SIAM J. Comput.*, 38(3):1175–1206, 2008.

[Mul12] Ketan Mulmuley. The GCT program toward the P vs. NP problem. *Commun. ACM*, 55(6):98–107, 2012.

[Pee96] René Peeters. Orthogonal representations over finite fields and the chromatic number of graphs. *Combinatorica*, 16(3):417–431, 1996.

[RR97] Alexander A. Razborov and Steven Rudich. Natural proofs. *J. Comput. Syst. Sci.*, 55(1):24–35, 1997.

[Sag01] Bruce E. Sagan. *The symmetric group*, volume 203 of *Graduate Texts in Mathematics*. Springer-Verlag, New York, second edition, 2001. Representations, combinatorial algorithms, and symmetric functions.

[Shi16] Yarolav Shitov. How hard is the tensor rank? *CoRR*, abs/1611.01559, 2016.

[SS16] Marcus Schaefer and Daniel Stefankovic. The complexity of tensor rank. *CoRR*, abs/1612.04338, 2016.

[SWZ17] Zhao Song, David P. Woodruff, and Peilin Zhong. Relative error tensor low rank approximation. *CoRR*, abs/1704.08246, 2017.

A. Algebraic complexity classes

In this section, we give some background on algebraic complexity theory. For a more comprehensive background, reader is referred to [Bür00]. We use $\mathbb{F}$ to denote the base field used in this section.

An arithmetic circuit over the field $\mathbb{F}$ is an acyclic finite directed graph, which has three kind of nodes, namely input nodes, gates, output nodes. The input nodes are labeled by variables from $\{X_1, X_2, \dots\}$ or by constants in $\mathbb{F}$. If all constants belong to $\{-1, 0, 1\}$, then the circuit is called constant-free. All nodes except the input nodes have fan-in 2 and are labeled by $+, -, \times$ or $/$, one calls these gates. The circuit is called division-free if there are no division nodes. The outputs nodes are the nodes which have no outgoing edges.

In general, an arithmetic circuit can have many output nodes but often one only considers arithmetic circuits with a single output node. In case there is exactly one output node, it is easy to see that the circuit computes a rational function in $\mathbb{F}(X_1, X_2, \dots)$ in the obvious way. The number of nodes in the circuit is called the size of the circuit.

Similar to circuit complexity of Boolean functions, one can study the arithmetic circuit complexity of polynomial families.

Definition 27. *The arithmetic circuit complexity $L(f)$ of a polynomial f is defined as the minimum size of an arithmetic circuit which computes f. The L^0 complexity $L^0(f)$ of an integer polynomial f is defined as the minimum size of a division-free and constant-free arithmetic circuit computing f.*

Now we can define the algebraic analogue of P .

Definition 28. *Complexity class VP is the set of polynomial families (f_n) such that both $\deg(f_n)$ and $L(f_n)$ are polynomially bounded functions of n.*

In this paper, we shall not be concerned with the complexity class VP but with complexity class VP^0. For this purpose, we define the following notion.

The formal degree of a node is inductively defined as follows: input nodes have formal degree 1. The formal degree of an addition or subtraction gate is the maximum of the formal degrees of the two incoming nodes, and the formal degree of a multiplication node is the gate of these formal degrees. The formal degree of a circuit is defined as the formal degree of its output node.

Definition 29. *A polynomial family (f_n) is in VP^0 iff there exists a sequence (C_n) of division-free and constant-free arithmetic circuits such that C_n computes f_n and the size and the formal degree of C_n are polynomially bounded functions of n. A polynomial family $(f_n(X_1, X_2, \dots, X_{p(n)}))$ is in VNP^0 iff there exists a sequence $(f_n(X_1, X_2, \dots, X_{q(n)}))$ in VP^0 such that*

$$(f_n(X_1, X_2, \dots, X_{p(n)})) = \sum_{e \in \{0,1\}^{q(n)-p(n)}} (f_n(X_1, X_2, \dots, X_{p(n)}, e_1, e_2, \dots, e_{q(n)-p(n)})).$$

It is know that permanent polynomial family $(\mathrm{Per}_n) = \sum_{\sigma \in S_n} X_{i,\sigma(i)}$ is VNP-complete under notion of p-projections [Bür00] if $\mathrm{char}(\mathbb{F}) \neq 2$. The polynomial family (Per_n) even belongs to VNP^0. Valiant's Hypothesis states that $(\mathrm{Per}_n) \notin \mathsf{VP}^0$ or even $(\mathrm{Per}_n) \notin \mathsf{VP}$.

B. Alternative definition of closures

A weaker notion of completion rank would be the following: We think of $A_1, \dots, A_m$ being fixed and only A_0 can be approximated. That is, we project $C_r^{m,n}$ down to $(A_0, \lambda_1, \dots, \lambda_m)$ and work with this variety, that is, all tuples such that $\mathrm{rk}(A_0 + \lambda_1 A_1 + \dots \lambda_m A_m) \leq r$. Then we project onto the first component, that is, the matrix, and take the limit there.

This version of completion rank might be closer to the actual application: We can replace the variables in the matrix by values such that we get a low rank approximation of the given matrix A_0. The version of Definition 9 is more general, we are even allowed to replace the variables "approximately" by modifying related entries, however, with a higher approximation order. The illustrate this effect, we study the following example:

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ x & y & 1 \end{pmatrix}.$$

Constants belong to the matrix A_0, i.e., A_0 is the identity matrix, the x variable represents A_1 and the y variable represents A_2. No matter how we substitute the variables, the upper-left 2×2-minor of any close-enough approximation to this matrix will be nonzero. (As a side remark, note that we can however bring down the rank to 2 by substituting $y \mapsto 1/\epsilon$.)

On the other hand, if we work with approximations to A_1 and A_2, we can bring down the border completion rank down to one. Let

$$\tilde{A}_1 := \begin{pmatrix} \epsilon & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} \quad \text{and} \quad \tilde{A}_2 := \begin{pmatrix} 0 & 0 & 0 \\ 0 & \epsilon & 0 \\ 0 & 1 & 0 \end{pmatrix}.$$

Then

$$\begin{pmatrix} 1 & 0 & 0 \\ 0 & 1 & 0 \\ 0 & 0 & 1 \end{pmatrix} - \frac{1}{\epsilon} \begin{pmatrix} \epsilon & 0 & 0 \\ 0 & 0 & 0 \\ 1 & 0 & 0 \end{pmatrix} - \frac{1}{\epsilon} \begin{pmatrix} 0 & 0 & 0 \\ 0 & \epsilon & 0 \\ 0 & 1 & 0 \end{pmatrix} = \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ -1/\epsilon & -1/\epsilon & 1 \end{pmatrix}$$

has rank 1, that is, $\underline{\mathrm{CR}}(A_0, A_1, A_2) = 1$.

This example shows that the two definitions actually differ. From the perspective of algebraic geometry and complexity, Definition 9 seems to be the right definition. Our hardness proofs and our barrier also works for the other definition, in fact, they are even easier.