A Polynomial Degree Bound on Equations for Non-rigid Matrices and Small Linear Circuits

Mrinal Kumar* Ben Lee Volk †

Abstract

We show that there is an equation of degree at most $\mathsf{poly}(n)$ for the (Zariski closure of the) set of the non-rigid matrices: that is, we show that for every large enough field $\mathbb{F}$, there is a non-zero n^2-variate polynomial $P \in \mathbb{F}(x_{1,1}, \ldots, x_{n,n})$ of degree at most $\mathsf{poly}(n)$ such that every matrix M which can be written as a sum of a matrix of rank at most $n/100$ and sparsity at most $n^2/100$ satisfies $P(M) = 0$. This confirms a conjecture of Gesmundo, Hauenstein, Ikenmeyer and Landsberg [GHIL16] and improves the best upper bound known for this problem down from $\exp(n^2)$ [KLPS14, GHIL16] to $\mathsf{poly}(n)$.

We also show a similar polynomial degree bound for the (Zariski closure of the) set of all matrices M such that the linear transformation represented by M can be computed by an algebraic circuit with at most $n^2/200$ edges (without any restriction on the depth). As far as we are aware, no such bound was known prior to this work when the depth of the circuits is unbounded.

Our methods are elementary and short and rely on a polynomial map of Shpilka and Volkovich [SV15] to construct low degree "universal" maps for non-rigid matrices and small linear circuits. Combining this construction with a simple dimension counting argument to show that any such polynomial map has a low degree annihilating polynomial completes the proof.

As a corollary, we show that any derandomization of the polynomial identity testing problem will imply new circuit lower bounds. A similar (but incomparable) theorem was proved by Kabanets and Impagliazzo [KI04].

1 Introduction

1.1 Equations for varities in algebraic complexity theory

Let $V \subseteq \mathbb{F}^n$ be a (not necessarily irreducible) affine variety and let $\mathbf{I}(V)$ denote its ideal. A non-zero polynomial $P \in \mathbf{I}(V)$ is called an *equation* for V. An equation for V may serve as a "proof" that a point $\mathbf{x} \in \mathbb{F}^n$ is *not* in V, by showing that $P(\mathbf{x}) \neq 0$.

A fundamental observation of the Geometric Complexity Theory program is that many important circuit lower bounds problems in algebraic complexity theory fit naturally into the setting of showing that a point $\mathbf{x}$ lies outside a variety V [MS01, BIL+19]. In this formulation, one considers V to be the closure of a class of polynomials of low complexity, and $\mathbf{x}$ is the coefficient vector of the candidate hard polynomial.

*Department of Computer Science & Engineering, IIT Bombay. Email: `mrinal@cse.iitb.ac.in`

†Center for the Mathematics of Information, California Institute of Technology, USA. Email: `benleevolk@gmail.com`

Let $\Delta(V) := \min_{0 \neq P \in \mathbf{I}(V)}\{\deg(P)\}$. The quantity $\Delta(V)$ can be thought of as a measure of complexity for the geometry of the variety V. The quantity $\Delta(V)$ is a very coarse complexity measure. A recent line of work regarding *algebraic natural proofs* [FSV18, GKSS17] suggests to study the arithmetic circuit complexity of equations for varieties V that correspond to polynomials with small circuit complexity. Having $\Delta(V)$ growing like a polynomial in n is a necessary (but not a sufficient) condition for a variety V to have an algebraic natural proof for non-containment.

1.2 Rigid matrices

A matrix M is (r, s)-rigid if M cannot be written as a sum $R + S$ where $\mathsf{rank}(R) \leq r$ and S contains at most s non-zero entries. Valiant [Val77] proved that if A is $(\varepsilon n, n^{1+\delta})$-rigid for some constants $\varepsilon, \delta > 0$ then A cannot be computed by arithmetic circuits of size $O(n)$ and depth $O(\log n)$, and posed the problem of *explicitly* constructing rigid matrices with these parameters, which is still open. It is easy to prove that most matrices have much stronger rigidity parameters: over algebraically closed fields a generic matrix is $(r, (n-r)^2)$-rigid for any target rank r.

Let $\mathbb{F}$ be an algebraically closed field. Let $A_{r,s} \subseteq \mathbb{F}^{n\times n}$ denote the set of matrices which are not (r, s)-rigid. Let $V_{r,s} = \overline{A_{r,s}}$ denote the Zariski closure of $A_{r,s}$. A geometric study of $V_{r,s}$ was initiated by Kumar, Lokam, Patankar and Sarma [KLPS14]. Among other results, they prove that for every $s < (n-r)^2$, $\Delta(V_{r,s}) \leq n^{4n^2}$. A slightly improved (but still exponential) upper bound was obtained by Gesmundo, Hauenstein, Ikenmeyer and Landsberg [GHIL16], who also conjectured that for some $\varepsilon, \delta > 0$, $\Delta(V_{\varepsilon n, n^{1+\delta}})$ grows like a polynomial function in n. The following theorem which we prove in this note confirms this conjecture.

Theorem 1.1. *Let $\varepsilon < 1/25$, and let $\mathbb{F}$ be a field of size at least n^2. For every large enough n, there exists a non-zero polynomial $Q \in \mathbb{F}[x_{1,1}, \ldots, x_{n,n}]$, of degree at most n^3, which is a non-trivial equation for matrices which are* not *$(\varepsilon n, \varepsilon n^2)$-rigid. That is, for every such matrix M, $Q(M) = 0$.*

In fact, the conjecture of [GHIL16] was slightly weaker: they conjectured that $\Delta(U)$ is polynomial in n for every irreducible component U of $V_{\varepsilon n, n^{1+\delta}}$. As shown by [KLPS14], the irreducible components are in one-to-one correspondence with subsets of $[n] \times [n]$ of size $n^{1+\delta}$ corresponding to possible supports of the sparse matrix S.

As we observe in Remark 2.3, it is somewhat simpler to show that each of these irreducible components has an equation with a polynomial degree bound. However, since the number of such irreducible components is exponentially large, it is not clear if there is a single equation for the whole variety which is of polynomially bounded degree. We do manage to reverse the order of quantifiers and prove such an upper bound in Theorem 1.1. This suggests that the set of non-rigid matrices is much less complex than what one may suspect given the results of [KLPS14, GHIL16].

1.3 Circuits for linear transformations

The original motivation for defining rigidity was in the context of proving lower bounds for algebraic circuits [Val77]. If $A \in \mathbb{F}^{n\times n}$ is an $(\varepsilon n, n^{1+\delta})$-rigid matrix, for any $\varepsilon, \delta > 0$, then the linear transformation represented by A cannot be computed by an algebraic circuit of depth $O(\log n)$ and size $O(n)$.

Every algebraic circuit computing a linear transformation is without loss of generality a *linear* circuit. A linear circuit is a directed acyclic graph that has n inputs labeled $X_1, \ldots, X_n$ and n output nodes. Each edge is labeled by a scalar $\alpha \in \mathbb{F}$. Each node computes a linear function in $X_1, \ldots, X_n$ defined inductively. An internal node u with children, $v_1, \ldots, v_k$, connected to it by edges labeled $\alpha_1, \ldots, \alpha_k$, computes the linear function $\sum_i \alpha_i \ell_{v_i}$, where ℓ_{v_i} is the linear function computed by v_i, $1 \leq i \leq k$. The size of the circuit is the number of edges in the circuit.

It is possible to use similar techniques to those used in the proof of Theorem 1.1 in order to prove a polynomial upper bound on an equation for a variety containing all matrices $A \in \mathbb{F}^{n \times n}$ whose corresponding linear transformation can be computed by an algebraic circuit of size at most $n^2/200$ (even without restriction on the depth). Note that this is nearly optimal as any such linear transformation can be computed by a circuit of size n^2. More formally, we show the following.

Theorem 1.2. *Let $\mathbb{F}$ be a field of size at least n^2. For every large enough n, there exists a non-zero polynomial $Q \in \mathbb{F}[x_{1,1}, \ldots, x_{n,n}]$, of degree at most n^3, which is a non-trivial equation for matrices which are computed by algebraic circuit of size at most $n^2/200$.*

Our proofs are based on a dimension counting arguments, and are therefore non-constructive and do not give explicit equations for the relevant varieties. It thus remains a very interesting open problem to provide explicit low-degree equations for any of the varieties considered in this paper. Here "explicit" means a polynomial which has arithmetic circuits of size $\mathsf{poly}(n)$.[1] The question of whether such equations exists has a win-win flavor: if they do, this can aid in explicit constructions of rigid matrices, and on the other hand, if all equations are hard, we have identified a family of polynomials which requires super-polynomial arithmetic circuits. Assuming the existence of a polynomial time algorithm for polynomial identity testing, we are able to make this connection formal.

Let PIT denote the set of strings which describe arithmetic circuits (say, over $\mathbb{C}$) which compute the zero polynomial. It is well known that $\mathsf{PIT} \in \mathsf{coRP}$. Kabanets and Impagliazzo [KI04] proved that certain circuit lower bounds follow from the assumption that $\mathsf{PIT} \in \mathsf{P}$. As a corollary to Theorem 1.2, we are able to prove theorem of a similar kind.

Corollary 1.3. *Suppose $\mathsf{PIT} \in \mathsf{P}$. Then at least one of the following is true:*

1. *There exists a family of n-variate polynomial of degree $\mathsf{poly}(n)$ over $\mathbb{C}$, which can be computed (as its list of coefficients, given the input 1^n) in PSPACE, which does not have polynomial size constant free arithmetic circuit.*
2. *there exists a family of matrices, constructible in polynomial time with an NP oracle (given the input 1^n), which requires linear circuits of size $\Omega(n^2)$.*

A *constant free arithmetic circuit* is an arithmetic circuit which is only allowed to use the constants $\{0, \pm 1\}$.

A different way to interpret Corollary 1.3 is as saying that at least one of the following three lower bound results hold: either $\mathsf{PIT} \notin \mathsf{P}$, or (at least) one of the two circuit lower bounds stated in the corollary. We emphasize that the result holds under *any* (even so-called *white box*) derandomization of PIT.

Our statement is similar to, but incomparable with the result of Kabanets and Impagliazzo [KI04] who proved that if $\mathsf{PIT} \in \mathsf{P}$ then either the permanent does not have polynomial size constant free arithmetic circuits, or $\mathsf{NEXP} \not\subseteq \mathsf{P/poly}$.

Since $(\varepsilon n, \varepsilon n^2)$-rigid matrices have linear circuit of size $3\varepsilon n^2$, the last item of Corollary 1.3 in particular implies a conditional construction of $(\Omega(n), \Omega(n^2))$-rigid matrices (it is also possible to directly use Theorem 1.1 instead of Theorem 1.2 to deduce this result). Unconditional constructions of rigid matrices in polynomial time with an NP oracle were recently given in [AC19, BHPT20]. However, the rigidity parameters in these papers are not enough to imply circuit lower bounds (furthermore, even optimal rigidity parameters are not enough to imply $\Omega(n^2)$ lower bounds for general linear circuits).

Since it is widely believed that $\mathsf{PIT} \in \mathsf{P}$, the answer to which of the last two items of Corollary 1.3 holds boils down to the question of whether there exists an equation for non-rigid

[1] Although one may consider other, informal notions of explicitness which could nevertheless be helpful.

matrices of degree $\mathsf{poly}(n)$ and circuit size $\mathsf{poly}(n)$. If determining if a matrix is rigid is coNP-hard (as is known for some restricted ranges of parameters [MS10]), it is tempting to also believe that the equations should not be easily computable, as they provide "proof" for rigidity which can be verified in randomized polynomial time. However, it could still be the case that those equations that have polynomial size circuits only prove the rigidity of "easy" instances.

2 Degree Upper Bound for Non-Rigid Matrices

In this section, we prove Theorem 1.1. A key component of the proof is the use of the following construction, due to Shpilka and Volkovich, which provides an explicit low-degree polynomial map on a small number of variables, which contains all sparse matrices in its image. For completeness, we provide the construction and prove its basic property.

Lemma 2.1 ([SV15])**.** *Let $\mathbb{F}$ be a field such that $|\mathbb{F}| > n$. Then for all $k \in \mathbb{N}$, there exists an explicit polynomial map $\mathrm{SV}_{n,k}(\mathbf{x}, \mathbf{y}) : \mathbb{F}^{2k} \to \mathbb{F}^n$ of degree at most n such that for any subset $T = \{i_1, \ldots, i_k\} \subseteq [n]$ of size k, there exists a setting $\mathbf{y} = \boldsymbol{\alpha}$ such that $\mathrm{SV}(\mathbf{x}, \boldsymbol{\alpha})$ is identically zero on every coordinate $j \notin T$, and equals x_j in coordinate i_j for all $j \in [k]$.*

Proof. Arbitrarily pick distinct $\alpha_1, \ldots \alpha_n \in \mathbb{F}$, and let $u_1, \ldots, u_n$ be their corresponding Lagrange's interpolation polynomials, i.e., polynomials of degree at most $n-1$ such that $u_i(\alpha_j) = 1$ if $j = i$ and 0 otherwise (more explicitly, $u_i(z) = \frac{\prod_{j \neq i}(z-\alpha_j)}{\prod_{j \neq i}(\alpha_i - \alpha_j)}$).

Let $P_i(x_1, \ldots, x_k, y_1, \ldots, y_k) = \sum_{j=1}^{k} u_i(y_j) \cdot x_j$, and finally let

$$\mathrm{SV}_{n,k}(\mathbf{x}, \mathbf{y}) = (P_1(\mathbf{x}, \mathbf{y}), \ldots, P_n(\mathbf{x}, \mathbf{y})).$$

It readily follows that given $T = \{i_1, \ldots, i_k\}$ as in the statement of the lemma, we can set $y_j = \alpha_{i_j}$ for $j \in [k]$ to derive the desired conclusion. The upper bound on the degree follows by inspection. ☐

As a step toward the proof of Theorem 1.1, we show there is a polynomial map on much fewer than n^2 variables with degree polynomially bounded in n such that its image contains every non-rigid matrix. In the next step, we show that the image of every such polynomial map has an equation of degree $\mathsf{poly}(n)$.

Lemma 2.2. *There exists an explicit polynomial map $P : \mathbb{F}^{4\varepsilon n^2} \to \mathbb{F}^{n \times n}$, of degree at most n^2, such that every matrix M which is not $(\varepsilon n, \varepsilon n^2)$ rigid lies in its image.*

Proof. Let $k = \varepsilon n^2$ and let $\mathbf{u}, \mathbf{v}, \mathbf{x}, \mathbf{y}$ denote disjoint tuples of k variables each.

Let U be a symbolic $n \times \varepsilon n$ matrix whose entries are labeled by the variables $\mathbf{u}$, and similarly let V be a symbolic $\varepsilon n \times n$ matrix labeled by $\mathbf{v}$. Let $\mathrm{UV}(\mathbf{u}, \mathbf{v}) : \mathbb{F}^{2k} \to \mathbb{F}^{n \times n}$ be the degree 2 polynomial map defined by the matrix multiplication UV.

Finally, let $P : \mathbb{F}^{4k} \to \mathbb{F}^{n \times n}$ be defined as

$$P(\mathbf{u}, \mathbf{v}, \mathbf{x}, \mathbf{y}) = \mathrm{UV}(\mathbf{u}, \mathbf{v}) + \mathrm{SV}_{n^2,k}(\mathbf{x}, \mathbf{y}),$$

where $\mathrm{SV}_{n^2,k}$ is as defined in Lemma 2.1.

Suppose now M is a non-rigid matrix, i.e., $M = R + S$ for R of rank εn as S which is εn^2-sparse. Decompose $R = U_0 V_0$ for $n \times \varepsilon n$ matrix U_0 and εn matrix V_0. Let T denote the support of S. For convenience we may assume $|T| = k$ (otherwise, pad with zeros arbitrarily). Let $\boldsymbol{\alpha} \in \mathbb{F}^k$ denote the setting for $\mathbf{y}$ in $\mathrm{SV}_{n^2,k}$ which maps $x_1, \ldots, x_k$ to T, and let $\mathbf{s} = (s_1, \ldots, s_k)$ denote the non-zero entries of S. Then

$$P(U_0, V_0, \mathbf{s}, \boldsymbol{\alpha}) = U_0 V_0 + S = R + S = M.$$ ☐

To complete the proof of Theorem 1.1, we now argue that the image of any polynomial map with parameters as in Lemma 2.2 has an equation of degree at most n^3.

Proof of Theorem 1.1. Let V_1 denote the subspace of polynomials over $\mathbb{F}$ in n^2 variables of degree at most n^3. Let V_2 denote the subspace of polynomials over $\mathbb{F}$ in $4\varepsilon n^2$ variables of degree at most n^5. Let P be as in Lemma 2.2, and consider the linear transformation $T : V_1 \to V_2$ given by $Q \mapsto Q \circ P$ (indeed, observe that since $\deg(Q) \leq n^3$ and $\deg(P) \leq n^2$, it follows that $\deg Q \circ P \leq n^5$).

We have that $\dim(V_1) = \binom{n^3+n^2}{n^2} \geq n^{n^2}$, whereas $\dim(V_2) = \binom{4\varepsilon n^2+n^5}{4\varepsilon n^2} \leq (2n^5)^{4\varepsilon n^2} < \dim(V_1)$ by the choice of ε, so that there exists a non-zero polynomial in the kernel of T, that is, $0 \neq Q_0 \in V_1$ such that $Q_0 \circ P \equiv 0$.

It remains to be shown that for any non-rigid matrix M, $Q_0(M) = 0$. Indeed, let M be a non-rigid matrix. By Lemma 2.2, there exist $\boldsymbol{\beta} \in \mathbb{F}^{4\varepsilon n^2}$ such that $P(\boldsymbol{\beta}) = M$. Thus, $Q_0(M) = Q_0(P(\boldsymbol{\beta})) = Q_0 \circ P(\beta) = 0$, as $Q_0 \circ P \equiv 0$. $\square$

Remark 2.3. *If the support of the sparse matrix is fixed a-priori to some set $S \subseteq [n] \times [n]$ of cardinality at most ϵn^2, then it is easier to come up with a universal map $\tilde{P}$ from $\mathbb{F}^{3\epsilon n^2} \mapsto \mathbb{F}^{n \times n}$ such that every matrix M whose rank can be reduced to at most ϵn by changing entries in the set S is contained in the image of $\tilde{P}$. Just consider $\tilde{P}(\mathbf{w}, \mathbf{x}, \mathbf{y}) = \mathrm{UV}(\mathbf{u}, \mathbf{v}) + W$, where W is a matrix such that for all $(i, j) \in [n] \times [n]$, if $(i, j) \in S$, then $W(i, j) = w_{i,j}$ and $W(i, j)$ is zero otherwise. Here, each $w_{i,j}$ is a distinct formal variable. Combined with the dimension comparison argument we used in the proof of Theorem 1.1, it can be seen that there is a non-zero low degree polynomial $\tilde{Q}$ such that $\tilde{Q} \circ \tilde{P} \equiv 0$. This argument provides a (different) equation of polynomial degree for each irreducible component of the variety of non-rigid matrices.* $\diamondsuit$

Remark 2.4. *It is possible to use the equation given in Theorem 1.1, and using the methods of [KLPS14], to construct "semi-explicit" $(\varepsilon n, \varepsilon n^2)$-rigid matrices. These are matrices whose entries are algebraic numbers (over $\mathbb{Q}$) with short description, which are non-explicit from the computational complexity point of view. However, such constructions are also known using different methods (see Section 2.4 of [Lok09]).* $\diamondsuit$

3 Degree Upper bound for Matrices with a Small Circuit

In this section, we prove Theorem 1.2. Our strategy, as before, is to observe that all matrices with a small circuit lie in the image of a polynomial map P on a small number of variables and small degree. Circuits of size s can have many different topologies and thus we first construct a "universal" linear circuit, of size $s' \leq s^4$, that contains as subcircuits all linear circuits of size s. Importantly, s' will affect the degree of P but not its number of variables. We note that this construction of universal circuits is slightly different from similar constructions in earlier work, e.g., in [Raz10]; the key difference being that a naive use of ideas in [Raz10] to obtain the map P seems to incur an asymptotic increase in the number of variables of P, which is unacceptable in our current setting.

Let $s \geq n$. We first define a universal graph G for size s. G has a set V_0 of n input nodes labeled $X_1, \ldots X_n$ and a set V_{s+1} of n designated output nodes. In addition, G is composed of s disjoint sets of vertices $V_1, \ldots, V_s$, each contains s vertices.

Each vertex $v \in V_i$, for $0 \leq i \leq s+1$, has as its children all vertices $u \in V_j$ for all $0 \leq j < i$. It is clear than every directed acyclic graph with s edges (and hence at most s vertices, and depth at most s) can be (perhaps non-uniquely) embedded in G as a subgraph.

We now describe the edge labeling. Let $s' \leq s^4$ be the number of edges in V and let e_i denote the i-th edge, $1 \leq i \leq s'$. The edge e_i is labeled by the i-th coordinate of the map $\mathrm{SV}_{s',s}(\mathbf{x}, \mathbf{y})$ given in Lemma 2.1.

Thus, the graph G with this labeling computes a linear transformation (over the field $\mathbb{F}(\mathbf{x}, \mathbf{y})$) in the variables $X_1, \ldots, X_n$. More explicitly, the (i, j)-th entry of the matrix $U(\mathbf{x}, \mathbf{y})$ representing this linear transformation is given by the sum, over all paths from X_i to the j-th output node, of the product of the edge labels on that path. This entry is a polynomial in $\mathbf{x}, \mathbf{y}$, so that we can think of U as a polynomial map from $\mathbb{F}^{2s}$ to $\mathbb{F}^{n^2}$.

Lemma 3.1. *The map $U(\mathbf{x}, \mathbf{y})$ defined above contains in its image all $n \times n$ matrices A whose corresponding linear transformation can be computed by a linear circuit of size at most s. The degree of U is at most $s' \cdot (s + 1)$.*

Proof. Let A be a matrix whose linear transformation is computed by a size s circuit C. The graph of C can be embedded as a subgraph in the graph G constructed above (if the embedding is not unique, pick one arbitrarily). Let $e_{i_1}, \ldots, e_{i_s}$ be the edges of this subgraph, and let $\boldsymbol{\beta} = (\beta_1, \ldots, \beta_s)$ be their corresponding labels in C. By the properties of the map $\mathrm{SV}_{s',s}(\mathbf{x}, \mathbf{y})$ given in Lemma 2.1, it is possible to set the tuple of variables $\mathbf{y}$ to field elements $\alpha_1, \ldots, \alpha_s$ such that the j-th coordinate of $\mathrm{SV}(\boldsymbol{\beta}, \boldsymbol{\alpha})$ equals β_i if $j = i_k$ for some $1 \leq k \leq s$ the 0 otherwise. Observe that under this labeling of the edges, the circuit G computes the same transformation as the circuit C. Hence $U(\boldsymbol{\beta}, \boldsymbol{\alpha}) = A$.

To upper bound the degree of U, note that each edge label in G is a polynomial of degree s', and each path is of length at most $s + 1$. □

Analogous to the proof of Theorem 1.1, we now observe via a dimension counting argument that the image of the polynomial map $U(\mathbf{x}, \mathbf{y})$ has a equation of degree at most n^3. This would complete the proof of Theorem 1.2.

Proof of Theorem 1.2. As before, let V_1 denote the subspace of polynomials over $\mathbb{F}$ in n^2 variables of degree at most n^3. Let V_2 denote the subspace of polynomials over $\mathbb{F}$ in $n^2/100$ variables of degree at most n^{30}. Let U be the map given by Lemma 3.1 for $s = n^2/200$ so that $s' \leq n^8$, and the degree of U is at most $s'(s+1) \leq n^{10}$. Now, consider the linear transformation $T : V_1 \to V_2$ given by $Q \mapsto Q \circ U$.

Once again, we compute that $\dim(V_1) = \binom{n^3+n^2}{n^2} \geq n^{n^2}$, whereas $\dim(V_2) = \binom{n^2/100+n^{30}}{n^2/100} \leq (2n^{30})^{n^2/100} < \dim(V_1)$, so that there exists a non-zero polynomial in the kernel of T, that is, $0 \neq Q_0 \in V_1$ such that $Q_0 \circ U \equiv 0$.

By Lemma 3.1, if A has a circuit of size $n^2/200$, it is in the image of U, so that $Q_0(A) = 0$. □

4 Degree upper bound for three dimensional tensors

Another algebraic object which is closely related to proving circuit lower bounds is the set of three dimensional tensors of high rank. A three dimensional tensor of rank at least r implies a lower bound of r on an arithmetic circuit computing the bi-linear function associated with the tensor. Our arguments also provide polynomial degree upper bounds for the set of tensors of (border) rank at most $n^2/300$.

Lemma 4.1. *Let $\mathbb{F}$ be any field. There is a polynomial map $P : \mathbb{F}^{n^2/100} \to \mathbb{F}^{n^3}$ of degree at most 3 such that for every 3 dimensional tensor $\tau : [n]^3 \to \mathbb{F}$ of rank at most $n^2/300$ lies in its image.*

Proof. This follows immediately from the definition.

Indeed, let $r = n^2/300$. Let $\mathbf{u}_1, \ldots, \mathbf{u}_r, \mathbf{v}_1, \ldots, \mathbf{v}_r, \mathbf{w}_1, \ldots, \mathbf{w}_r$ be disjoint n tuples of variables. Let U be a tensor of rank at most r over the ring $\mathbb{F}[\mathbf{u}_1, \ldots, \mathbf{u}_r, \mathbf{v}_1, \ldots, \mathbf{v}_r, \mathbf{w}_1, \ldots, \mathbf{w}_r]$ defined as follows.

$$U(\mathbf{u}, \mathbf{v}, \mathbf{w}) = \sum_{i=1}^{r} \mathbf{u}_i \otimes \mathbf{v}_i \otimes \mathbf{w}_i \, .$$

From the definition of U, it can be readily observed that for every tensor $\tau : \mathbb{F}^{[n]^3} \to \mathbb{F}$ of rank at most r, there is a setting $\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}$ of the variables in $\mathbf{u}, \mathbf{v}, \mathbf{w}$ respectively such that $U(\boldsymbol{\alpha}, \boldsymbol{\beta}, \boldsymbol{\gamma}) = \tau$. Moreover, each of the coordinates of U is a polynomial of degree equal to three in the variables in $\mathbf{u}, \mathbf{v}, \mathbf{w}$. Let P be the degree three polynomial map which maps the variables $\mathbf{u}_1, \ldots, \mathbf{u}_r, \mathbf{v}_1, \ldots, \mathbf{v}_r$ and $\mathbf{w}_1, \ldots, \mathbf{w}_r$ to the coordinates of U. $\square$

We now argue that for every polynomial map P given by Lemma 4.1 has an equation of not too large degree.

Theorem 4.2. *Let $\mathbb{F}$ be any field. There exists a non-zero polynomial $Q \in \mathbb{F}[x_{1,1,1}, \ldots, x_{n,n,n}]$, of degree at most n^4, which is a non-trivial equation for three dimensional tensors $\tau : [n] \times [n] \times [n] \mapsto \mathbb{F}$ of rank at most $n^2/300$.*

Proof. As before, let V_1 denote the subspace of polynomials over $\mathbb{F}$ in n^3 variables of degree at most n^4 and let V_2 denote the subspace of polynomials over $\mathbb{F}$ in $n^3/100$ variables of degree at most $3n^4$. Let P be the map given by Lemma 4.1. Now, consider the linear transformation $T : V_1 \to V_2$ given by $Q \mapsto Q \circ P$.

Observe that $\dim(V_1) = \binom{n^4+n^3}{n^3} \geq n^{n^3}$, whereas $\dim(V_2) = \binom{n^3/100+3n^4}{n^3/100} \leq (2n^4)^{n^3/100} < \dim(V_1)$, so that there exists a non-zero polynomial in the kernel of T, that is, $0 \neq Q_0 \in V_1$ such that $Q_0 \circ P \equiv 0$.

By Lemma 4.1, if τ is a tensor of rank at most $n^2/300$, then it is in the image of P, and thus $Q_0(\tau) = 0$. $\square$

The arguments here also generalize to tensors in higher dimensions. In particular, the following analog of Lemma 4.1 is true.

Lemma 4.3. *Let $\mathbb{F}$ be any field. Then, for all $n, d \in \mathbb{N}$, there is a polynomial map $P : \mathbb{F}^{n^{d-1}/100} \to \mathbb{F}^{n^d}$ of degree at most d such that for every d dimensional tensor $\tau : [n]^{\otimes d} \to \mathbb{F}$ of rank at most $n^{d-1}/100d$ lies in its image.*

Combining this lemma with a dimension comparison argument analogous to that in the proof of Theorem 4.2 gives the following theorem. We skip the details of the proof.

Theorem 4.4. *For every field $\mathbb{F}$ and for all $n, d \in \mathbb{N}$, there exists a non-zero polynomial Q on n^d variables and degree at most n^{2d}, which is a non-trivial equation for d dimensional tensors $\tau : [n]^{\otimes d} \to \mathbb{F}$ of rank at most $n^{d-1}/100d$.*

We remark that a similar methods can be used to prove the existence of an equation of degree $\mathsf{poly}(n)$ for three dimensional tensors of *slice rank* (see, e.g., [BIL+19]) at most, say, $n/1000$. The existence of such an equations was proved (using different techniques) in [BIL+19].

5 Applications to circuit lower bounds

In this section we prove Corollary 1.3. The strategy of the proof is simple: the proof of Theorem 1.2 implies a PSPACE algorithm which produces a sequence of polynomials which are equations for the set of matrices with small linear circuits. If those equations require large circuits, we are done, and if not, then there exists an equation with small circuits which (assuming $\mathsf{PIT} \in \mathsf{P}$) can be found using an NP-oracle. Using, once again, the assumption that $\mathsf{PIT} \in \mathsf{P}$, we can also find deterministically a matrix on which the equation evaluates to non-zero, which implies the matrix requires large linear circuits.

There are some technical difficulties involved with this plan which we now describe. The first problem is that even arithmetic circuits of small size can have large description as bit strings, due to the field constants appearing in the circuits. To prevent this issue, we only consider *constant free* arithmetic circuits, which are only allowed inputs labeled by $\{0, \pm 1\}$ (but can still compute other constants in the circuit using arithmetic operations).

The second problem is that, in order to be able to find a non-zero of the equation in the last step of the algorithm (using the mere assumption that $\mathsf{PIT} \in \mathsf{P}$), we need not only the size of the circuit but also its *degree* to be bounded by $\mathsf{poly}(n)$. Of course, by Theorem 1.2 the exists such a circuit, but we need to be able to prevent a malicious prover from providing us with a $\mathsf{poly}(n)$ size circuit of exponential degree, and it is not known how to compute the degree of a circuit in deterministic polynomial time, even assuming $\mathsf{PIT} \in \mathsf{P}$. To solve this issue, we use an idea of Malod and Portier [MP08], who showed that any polynomial with circuit of size $\mathsf{poly}(n)$ and degree d also has a *multiplicatively disjoint* (MD) circuit of size $\mathsf{poly}(n, d)$. An MD circuit is a circuit in which any multiplication gates multiplies two disjoint subcircuits. This is a syntactic notion which is easy to verify efficiently and deterministically, and an MD circuit of size s is guaranteed to compute a polynomial of degree at most s.

A final technical issue is that the notion of MD circuits does not fit perfectly within the framework of constant free circuits. Therefore we use the notion of "almost MD" circuits, which allow for the case which the inputs to a multiplciation gates are not disjoint, as long as at least one of them is the root of a subcircuit in which only constants appear.

Definition 5.1. *We say a gate v in a circuit is* constant producing *(CP) if in the subcircuit rooted at v, all input nodes are field constants.*

An almost-MD circuit *is a circuit where every multiplication gate either multiplies two disjoint subcircuits, or at least one of its children is constant producing.* $\Diamond$

Lemma 5.2. *Suppose f is an n-variate polynomial of degree $\mathsf{poly}(n)$ which has a constant free arithmetic circuit of degree $\mathsf{poly}(n)$. Then f has a constant free almost-MD circuit of size $\mathsf{poly}(n)$.*

Proof. Let C_0 be a constant free arithmetic circuit for f. We first homogenize the circuit C_0 to obtain a circuit C_1 (a homogeneous circuit is a circuit in which every gate computes a homogeneous polynomial [SY10]). Since C_1 is homogeneous, all the gates which compute non-zero field constants are CP gates. We then eliminate all gates which compute constants by allowing the edges entering sum gates to be labeled by field scalars, and interpreting a sum gate as computing a linear combination whose coefficients are given by the edge labels. We call this circuit C_2. This step does not maintain constant-freeness. However, every label appearing on the edges of C_2 was computed in C_1, so it can be computed by a constant-free arithmetic circuit of polynomial size.

We now do the transformation detailed in [MP08] to C_2 to obtain an MD circuit C_3, which has labels on the edges. This step does not produce new constants. Finally, we convert C_3 to an almost-MD constant free circuit C_4, by re-computing every label appearing on the edge using a fresh subcircuit for each label, and rewiring the circuit (which will convert the circuit from an MD circuit to an almost MD circuit). These subcircuits are guaranteed to have polynomial size constant free circuits since these constant were all computed in C_0, which keeps the total size $\mathsf{poly}(n)$. $\square$

For circuits which compute low-degree polynomials, the mere existence of an algorithm for the decision version of PIT allows one to construct an algorithm for the search version.

Lemma 5.3. *Suppose $\mathsf{PIT} \in \mathsf{P}$. Then there is a polynomial time algorithm that given a non-zero almost-MD arithmetic circuit C of size s computing an n-variate polynomial, finds in time $\mathsf{poly}(n, s)$ an element $\mathbf{a} \in \mathbb{C}^n$ such that $C(\mathbf{a}) \neq 0$.*

Proof. We abuse notation by denoting by C also the polynomial computed by the circuit C. Note that since C is almost-MD, the degree of C is at most s. Thus, there exists $a_1 \in \{0, 1, \ldots, s\}$ such that $C(a_1, x_2, \ldots, x_n)$ is a non-zero polynomial in $x_2, \ldots, x_n$. By iterating over those $s + 1$

values from 0 to s and using the assumption that $\mathsf{PIT} \in \mathsf{P}$, we can find such a value for a_1 in time $\mathsf{poly}(n, s)$. We then continue in the same manner with the rest of the variables. $\square$

As we noted above, the assumption that C is almost-MD was used in Lemma 5.3 to bound the degree of the circuit. It is also useful because it is easy to decide in deterministic polynomial time whether a circuit is almost-MD. We now complete the proof of Corollary 1.3.

Proof of Corollary 1.3. For every n, the proof of Theorem 1.2 provides an equation Q_n for the set of $n \times n$ matrices with small linear circuits. This polynomial can be found by solving a linear system of equations in a linear space whose dimension is $\exp(\mathsf{poly}(n))$. Using standard, small space algorithm for linear algebra [BvzGH82, ABO99], this implies that there exists a fixed PSPACE algorithm which, on input 1^n, outputs the list of coefficients of the polynomial Q_n.

Consider now the family $\{Q_n\}_{n \in \mathbb{N}}$. If for any constant $k \in \mathbb{N}$ there exist infinitely many $n \in \mathbb{N}$ such that Q_n requires circuits of size at least n^k, it follows (by definition) that the PSPACE algorithm above outputs a family of polynomials with super-polynomial constant-free arithmetic circuits.

We are thus left to consider the case that there exists a constant $k \in \mathbb{N}$ such that for all large enough $n \in \mathbb{N}$, Q_n can be computed by circuits of size n^k. By Lemma 5.2, we may assume without loss of generality that these circuits are almost-MD circuits. Further suppose $\mathsf{PIT} \in \mathsf{P}$. We will show how to construct a matrix in polynomial time with an NP oracle which requires large linear circuits.

Consider the language L of pairs $(1^n, x)$ such that there exists a string y of length at most n^k such that xy describes an almost-MD circuit C such that C is non-zero, and $C \circ U \equiv 0$, where U is the polynomial map given in the proof of Theorem 1.2.

Assuming $\mathsf{PIT} \in \mathsf{P}$, the language L is in NP, and by assumption for every large enough n there exists such a circuit. Thus, we can use the NP oracle to construct such a circuit C bit by bit. Finally, using Lemma 5.3 we can output a matrix M such that $C(M) \neq 0$.

By the properties of the circuit C and the map U, M does not have linear circuits of size less than $n^2/200$. $\square$

Many variations of Corollary 1.3 can be proved as well, with virtually the same proof. By slightly modifying the language L used in the proof, it is possible to prove the same result even under the assumption $\mathsf{PIT} \in \mathsf{NP}$ (recall that $\mathsf{PIT} \in \mathsf{coRP}$). A similar statements also holds over finite fields of size $\mathsf{poly}(n)$, in which case the proof is simpler since there are no issues related to the bit complexity of the first constants. Finally, an analog of Corollary 1.3 also holds for tensor rank, by using Theorem 4.2 instead of Theorem 1.2: that is, assuming $\mathsf{PIT} \in \mathsf{P}$, either there exists a construction of a hard polynomial in PSPACE, or an efficient construction with an NP oracle of a 3-dimensional tensor of rank $\Omega(n^2)$. We remark that for tensors of large rank there are no analogs of [AC19, BHPT20], i.e., there do not exist even constructions with an NP oracle of tensors with slightly super-linear rank.

References

[ABO99] Eric Allender, Robert Beals, and Mitsunori Ogihara. The Complexity of Matrix Rank and Feasible Systems of Linear Equations. *Comput. Complex.*, 8(2):99–126, 1999.

[AC19] Josh Alman and Lijie Chen. Efficient Construction of Rigid Matrices Using an NP Oracle. In *Proceedings of the 60th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2019)*, pages 1034–1055. IEEE Computer Society, 2019.

[BHPT20] Amey Bhangale, Prahladh Harsha, Orr Paradise, and Avishay Tal. Rigid Matrices From Rectangular PCPs. *Electronic Colloquium on Computational Complexity (ECCC)*, 27:75, 2020.

[BIL+19] Markus Bläser, Christian Ikenmeyer, Vladimir Lysikov, Anurag Pandey, and Frank-Olaf Schreyer. Variety Membership Testing, Algebraic Natural Proofs, and Geometric Complexity Theory. *CoRR*, abs/1911.02534, 2019. Pre-print available at `arXiv:1911.02534`.

[BvzGH82] Allan Borodin, Joachim von zur Gathen, and John E. Hopcroft. Fast Parallel Matrix and GCD Computations. *Inf. Control.*, 52(3):241–256, 1982.

[FSV18] Michael A. Forbes, Amir Shpilka, and Ben Lee Volk. Succinct Hitting Sets and Barriers to Proving Lower Bounds for Algebraic Circuits. *Theory of Computing*, 14(1):1–45, 2018.

[GHIL16] Fulvio Gesmundo, Jonathan D. Hauenstein, Christian Ikenmeyer, and J. M. Landsberg. Complexity of Linear Circuits and Geometry. *Foundations of Computational Mathematics*, 16(3):599–635, 2016.

[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. Pre-print available at `arXiv:1701.01717`.

[KI04] Valentine Kabanets and Russell Impagliazzo. Derandomizing Polynomial Identity Tests Means Proving Circuit Lower Bounds. *Computational Complexity*, 13(1-2):1–46, 2004. Preliminary version in the *35th Annual ACM Symposium on Theory of Computing (STOC 2003)*.

[KLPS14] Abhinav Kumar, Satyanarayana V. Lokam, Vijay M. Patankar, and Jayalal Sarma. Using Elimination Theory to Construct Rigid Matrices. *Computational Complexity*, 23(4):531–563, 2014.

[Lok09] Satyanarayana V. Lokam. Complexity Lower Bounds using Linear Algebra. *Foundations and Trends in Theoretical Computer Science*, 4(1-2):1–155, 2009.

[MP08] Guillaume Malod and Natacha Portier. Characterizing Valiant's algebraic complexity classes. *J. Complex.*, 24(1):16–38, 2008.

[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.

[MS10] Meena Mahajan and Jayalal Sarma. On the Complexity of Matrix Rank and Rigidity. *Theory Comput. Syst.*, 46(1):9–26, 2010.

[Raz10] Ran Raz. Elusive Functions and Lower Bounds for Arithmetic Circuits. *Theory of Computing*, 6(7):135–177, 2010.

[SV15] Amir Shpilka and Ilya Volkovich. Read-once polynomial identity testing. *Computational Complexity*, 24(3):477–532, 2015. Preliminary version in the *40th Annual ACM Symposium on Theory of Computing (STOC 2008)*.

[SY10] Amir Shpilka and Amir Yehudayoff. Arithmetic Circuits: A survey of recent results and open questions. *Foundations and Trends in Theoretical Computer Science*, 5:207–388, March 2010.

[Val77] Leslie G. Valiant. Graph-Theoretic Arguments in Low-Level Complexity. In *Proceedings of the 2nd International Symposium on the Mathematical Foundations of Computer Science (MFCS 1977)*, volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer, 1977.