Using Elimination Theory to construct Rigid Matrices

Abhinav Kumar[*]**, Satyanarayana V. Lokam**[†]**,
Vijay M. Patankar**[†]**, Jayalal Sarma M. N.**[‡]

ABSTRACT. The rigidity of a matrix A for target rank r is the minimum number of entries of A that must be changed to ensure that the rank of the altered matrix is at most r. Since its introduction by Valiant [22], rigidity and similar rank-robustness functions of matrices have found numerous applications in circuit complexity, communication complexity, and learning complexity. Almost all $n \times n$ matrices over an infinite field have a rigidity of $(n - r)^2$. It is a long-standing open question to construct infinite families of *explicit* matrices even with superlinear rigidity when $r = \Omega(n)$.

In this paper, we construct an infinite family of complex matrices with the largest possible, i.e., $(n - r)^2$, rigidity. The entries of an $n \times n$ matrix in this family are distinct primitive roots of unity of orders roughly $\exp(n^4 \log n)$. To the best of our knowledge, this is the first family of concrete (but not entirely explicit) matrices having maximal rigidity and a succinct algebraic description.

Our construction is based on elimination theory of polynomial ideals. In particular, we use results on the existence of polynomials in elimination ideals with effective degree upper bounds (effective Nullstellensatz). Using elementary algebraic geometry, we prove that the dimension of the affine variety of matrices of rigidity at most k is exactly $n^2 - (n - r)^2 + k$. Finally, we use elimination theory to examine whether the rigidity function is semicontinuous.

1 Introduction

Valiant [22] introduced the notion of matrix rigidity. The rigidity function $\mathrm{Rig}(A, r)$ of a matrix A for target rank r is defined to be the smallest number of entries of A that must be changed to ensure that the altered matrix has rank at most r. It is easy to see that for every $n \times n$ matrix A (over any field), $\mathrm{Rig}(A, r) \leqslant (n - r)^2$. Valiant also showed that, over an infinite field, almost all matrices have rigidity exactly $(n - r)^2$. It is a long-standing open question to construct infinite families of *explicit* matrices with superlinear rigidity for $r = \Omega(n)$. Here, by an explicit family, we mean that the $n \times n$ matrix in the family is computable by a deterministic Turing machine in time polynomial in n or by a Boolean circuit of size polynomial in n. Lower bounds on rigidity of explicit matrices are motivated by their numerous applications in complexity theory. In particular, Valiant showed that lower bounds of the form $\mathrm{Rig}(A, \epsilon n) = n^{1+\delta}$ (where ϵ and δ are some positive constants) imply that the linear transformation defined by A cannot be computed by arithmetic circuits of linear size and logarithmic depth consisting of gates that compute linear functions of their inputs. Since then, applications of lower bounds on rigidity and similar rank-robustness functions have been found in circuit complexity, communication complexity, and learning complexity ([7],

[*] abhinav@math.mit.edu, Department of Mathematics, MIT, USA.

[†] {satya,vij}@microsoft.com, Microsoft Research India, Bangalore, India.

[‡] jayalal@tsinghua.edu.cn, Institute for Theoretical Computer Science, Tsinghua University, Beijing, China.

[13], [15], [18], [19]). Two comprehensive surveys on this topic are [4] and [5]. Over finite fields, the best known lower bound for explicit A was first proved by Friedman [8] and is $\text{Rig}(A, r) = \Omega(\frac{n^2}{r} \log \frac{n}{r})$ for parity check matrices of good error-correcting codes. Over infinite fields, the same lower bound was proved by Shokrollahi, Spielman, and Stemann [21] for Cauchy matrices, Discrete Fourier Transform matrices of prime order (see [14]), and other families. Note that this type of lower bound reduces to the trivial $\text{Rig}(A, r) = \Omega(n)$ when $r = \Omega(n)$. In [16], lower bounds of the form $\text{Rig}(A, \epsilon n) = \Omega(n^2)$ were proved when $A = (\sqrt{p_{jk}})$ or when $A = (\exp(2\pi i / p_{jk}))$, where p_{jk} are the first n^2 primes. These matrices, however, are not explicit in the sense defined above.

In this paper, we construct an infinite family of complex matrices with the highest possible, i.e., $(n - r)^2$ rigidity. The entries of the $n \times n$ matrix in this family are primitive roots of unity of orders roughly $\exp(n^4 \log n)$. We show that the real parts of these matrices are also maximally rigid. Like the matrices in [16], this family of matrices is not explicit in the sense of efficient computability described earlier. However, one of the motivations for studying rigidity comes from algebraic complexity. In the world of algebraic complexity, any element of the ground field (in our case $\mathbb{C}$) is considered a primitive or atomic object. In this sense, the matrices we construct are explicitly described algebraic entities. To the best of our knowledge, this is the first construction giving an infinite family of non-generic/concrete matrices with maximum rigidity. It is still unsatisfactory, though, that the roots of unity in our matrices have orders exponential in n. Earlier constructions in [16] use roots of unity of orders $O(n^2)$ but the bounds on rigidity proved there are weaker: $n(n - cr)$ for some constant $c > 2$.

We pursue a general approach to studying rigidity based on elementary algebraic geometry and elimination theory. To set up the formalism of this approach, we begin by re-proving Valiant's result that the set of matrices of rigidity less than $(n - r)^2$ form a Zariski closed set in $\mathbb{C}^{n \times n}$, i.e., such matrices are solutions of a finite system of polynomial equations (hence a generic matrix has rigidity at least $(n - r)^2$). In fact, we prove a more general statement: the set of matrices of rigidity at most k has dimension (as an affine variety) exactly $n^2 - (n - r)^2 + k$. This sheds light on the geometric structure of rigid matrices. Our transversality argument in this context is clearer and cleaner than an earlier attempt in this direction (in the projective setting) by [11]. To look for specific matrices of high rigidity, we consider certain elimination ideals associated to matrices with rigidity at most k. A result in [1] using effective Nullstellensatz bounds [2], [9] shows that an elimination ideal of a polynomial ideal must always contain a nonzero polynomial with an explicit degree upper bound (Theorem 8). We then use simple facts from algebraic number theory to prove that a matrix whose entries are primitive roots of sufficiently high orders cannot satisfy any polynomial with such a degree upper bound. This gives us the claimed family of matrices of maximum rigidity.

Our primary objects of interest in this paper are the varieties of matrices with rigidity at most k. For a fixed k, we have a natural decomposition of this variety based on the patterns of changes. We prove that this natural decomposition is indeed a decomposition into *irreducible* components (Corollary 13). In fact, these components are defined by elimination ideals of determinantal ideals generated by all the $(r + 1) \times (r + 1)$ minors of an $n \times n$ matrix of indeterminates. Better effective upper bounds on the degree of a nonzero polynomial in

the elimination ideal of determinantal ideals than given by Theorem 8 would lead to similar improvements in the bound on the order of the primitive roots of unity we use to construct our rigid matrices. While determinantal ideals have been well-studied in mathematical literature, their elimination theory does not seem to have been as well-studied. Application to rigidity of these elimination ideals of determinantal ideals might be a natural motivation for studying them.

We next consider the question: given a matrix A, is there a small neighborhood of A within which the rigidity function is nondecreasing, i.e. such that every matrix in this neighborhood has rigidity at least equal to that of A? This is related to the notion of *semicontinuity* of the rigidity function. We give a family of examples to show that the rigidity function is in general not semicontinuous. However, the *specific* matrices we produce above, by their very construction, have neighborhoods within which rigidity is nondecreasing.

1.1   Definitions and Notations

Let F be a field. Then, by $M_n(F)$ we denote the algebra of $n \times n$ matrices over F. At times, when it is clear from the context, we will denote $M_n(F)$ by M_n. Let $X \in M_n(F)$. Then by X_{ij} we will denote the (i,j)-th entry of X. Given $X \in M_n(F)$, the support of X is defined as $Supp(X) := \{(i,j) \mid X_{ij} \neq 0 \in F\}$. Given a non-negative integer k, we define

$$S(k) := \{X \in M_n(F) : |Supp(X)| \leqslant k\}.$$

Thus, $S(k)$ is the set of matrices over F with at most k non-zero entries. A *pattern* π is a subset of the positions of an $n \times n$ matrix. Then, we define:

$$S(\pi) := \{X \in M_n(F) : Supp(X) \subseteq \pi\}.$$

Note that $S(k) = \cup_{|\pi|=k} S(\pi)$.

We say that a matrix X is (r,k)-rigid if changing at most k entries of X does not bring down the rank of the matrix to a value $\leqslant r$. More formally,

DEFINITION 1. *A matrix X is (r,k)-rigid if $rank(X+T) > r$ whenever $T \in S(k)$.*

DEFINITION 2. *The rigidity function $\mathrm{Rig}(X,r)$ is the smallest integer k for which the matrix X is not (r,k)-rigid. That is, $\mathrm{Rig}(X,r)$ is the minimum number of entries we need to change in the matrix X so that the rank becomes at most r:*

$$\mathrm{Rig}(X,r) := \min\{Supp(T) : \mathrm{rank}(X+T) \leqslant r\}.$$

Sometimes, we will allow T to be chosen in $M_n(L)$ for L an extension field of F. In this case we will denote the rigidity by $Rig(X,r,L)$.

Let $\mathrm{RIG}(n,r,k)$ denote the set of $n \times n$ matrices X such that $\mathrm{Rig}(X,r) = k$. Similarly, we define $\mathrm{RIG}(n,r,\geqslant k)$ to be the set of matrices of rigidity at least k and $\mathrm{RIG}(n,r,\leqslant k)$ to be the set of matrices of rigidity at most k. For a pattern π of size k, let $\mathrm{RIG}(n,r,\pi)$ be the set of matrices X such that for some $T_\pi \in S(\pi)$ we have $\mathrm{rank}(X+T_\pi) \leqslant r$. Then we have

$$\mathrm{RIG}(n,r,\leqslant k) = \bigcup_{\pi,|\pi|=k} \mathrm{RIG}(n,r,\pi).$$

1.2 Elimination Theory: Closure Theorem

We refer the reader to a standard text in algebraic geometry [6, 20] for the necessary background. Here we recall a basic result from Elimination Theory which is directly used in the paper. As the name suggests, Elimination Theory deals with elimination of a subset of variables from a given set of polynomial equations and finding the *reduced set* of polynomial equations (not involving the eliminated variables). The main results of Elimination Theory, especially the Closure Theorem, describe a precise relation between the reduced ideal and the given ideal, and its corresponding geometric interpretation.

Given an ideal $I = \langle f_1, \ldots, f_s \rangle \subseteq F[x_1, \ldots, x_n]$, the l-th *elimination ideal* I_l is the ideal of $F[x_{l+1}, \ldots, x_n]$ defined by $I_l := I \cap F[x_{l+1}, \ldots, x_n]$.

THEOREM 3.(Closure Theorem, page 125, Theorem 3 of [6])
Let I be an ideal of $F[x_1, \ldots, x_n, y_1, \ldots, y_m]$ and $I_n := I \cap F[y_1, \ldots, y_m]$ be the n-th elimination ideal associated to I. Let $V(I)$ and $V(I_n)$ be the subvarieties of $\mathbb{A}^{n+m}$ and $\mathbb{A}^m$ (the affine spaces over F of dimension $n + m$ and m respectively) defined by I and I_n respectively. Let p be the natural projection map from $\mathbb{A}^{n+m} \to \mathbb{A}^m$ (projection map onto the y-coordinates). Then,

1. *$V(I_n)$ is the smallest (closed) affine variety containing $p(V(I)) \subseteq \mathbb{A}^m$. In other words, $V(I_n)$ is the Zariski closure of $p(V(I))(\bar{F}) \subseteq \bar{F}^m$.*
2. *When $V(I)(\bar{F}) \neq \phi$, there is an affine variety W strictly contained in $V(I_n)$ such that $V(I_n) - W \subseteq p(V(I))$.*

2 Use of Elimination Theory

2.1 Determinantal Ideals and their Elimination Ideals

We would like to investigate the structure of the sets $\mathsf{RIG}(n, r, \leqslant k)$ and $\mathsf{RIG}(n, r, \pi)$ and their Zariski closures

$$
\begin{aligned}
\mathcal{W}(n, r, \leqslant k) &:= \overline{\mathsf{RIG}(n, r, \leqslant k)} \quad \text{and} \\
\mathcal{W}(n, r, \pi) &:= \overline{\mathsf{RIG}(n, r, \pi)}
\end{aligned}
$$

in the n^2-dimensional affine space of $n \times n$ matrices. Let X be an $n \times n$ matrix with entries being indeterminates $x_1, \ldots, x_{n^2}$. For a pattern π of k positions, let T_π be the $n \times n$ matrix with indeterminates $t_1, \ldots, t_k$ in the positions given by π. Note that saying $X + T_\pi$ has rank at most r is equivalent to saying that all its $(r + 1) \times (r + 1)$ minors vanish. Let us consider the ideal generated by these minors:

$$
I(n, r, \pi) := \left\langle Minors_{(r+1) \times (r+1)}(X + T_\pi) \right\rangle \subseteq F[x_1, \ldots, x_{n^2}, t_1, \ldots, t_k]. \tag{1}
$$

It then follows from the definition of rigidity that $\mathsf{RIG}(n, r, \pi)$ is the projection from $\mathbb{A}^{n^2} \times \mathbb{A}^k$ to $\mathbb{A}^{n^2}$ of the algebraic set $V(I(n, r, \pi))(F)$. Thus, if we define the elimination ideal

$$
EI(n, r, \pi) := I(n, r, \pi) \cap F[x_1, \ldots, x_{n^2}] \subseteq F[x_1, \ldots, x_{n^2}],
$$

then by the Closure Theorem (Theorem 3), we obtain

$$\mathcal{W}(n,r,\pi) = V(EI(n,r,\pi)). \tag{2}$$

Note that

$$\mathcal{W}(n,r,\leqslant k) = \bigcup_{\pi, |\pi|=k} \mathcal{W}(n,r,\pi).$$

2.2 Valiant's Theorem

The following theorem due to Valiant [22, Theorem 6.4, page 172] says that a generic matrix has rigidity $(n-r)^2$. That is, for $k < (n-r)^2$, the dimension of $\mathcal{W}(n,r,\leqslant k)$ is strictly less than n^2.

A reader familiar with Valiant's proof will realize that our proof is basically a rephrasing of Valiant's proof in the language of algebraic geometry. The point of this proof is to set up the formalism and use it later; in particular, when we compute the exact dimension of the rigidity variety $\mathcal{W}(n,r,\leqslant k)$.

THEOREM 4.(**Valiant**) *Let $n \geqslant 1, 0 < r < n$ and $0 \leqslant k < (n-r)^2$. Let $\mathcal{W} := \mathcal{W}(n,r,\leqslant k)$ be as above. Then,*

$$\dim(\mathcal{W}) < n^2.$$

PROOF.   Let $\pi \subseteq \{(i,j)|1 \leqslant i,j \leqslant n\}$ be a pattern of size k. Let τ be the index set of a fixed $r \times r$ minor. For a matrix B, let B_τ denote the minor of B indexed by τ. Define $\mathsf{RIG}(n,r,\pi,\tau)$ to be the set of all $n \times n$ matrices A that satisfy the following properties: there exists some $n \times n$ matrix T_π such that

1. $Supp(T_\pi) \subseteq \pi$,
2. $\text{rank}(A + T_\pi) = r$, and
3. $\det((A + T_\pi)_\tau) \neq 0$ where τ denotes the fixed $r \times r$ minor as above.

Recall that $S(\pi)$ is the set of matrices whose support is contained in π. Let us also define

$$\mathsf{RANK}(n,r,\tau) := \{C \in M_n \mid \text{rank}(C) = r \text{ and } \det(C_\tau) \neq 0\}.$$

By definition, every element $A \in \mathsf{RIG}(n,r,\pi,\tau)$ can be written as $C - T_\pi$, with $C \in \mathsf{RANK}(n,r,\tau)$ and $T_\pi \in S(\pi)$.

We state the following lemma without proof. (Details can be found in the full version [10]).

LEMMA 5. $\dim(\mathsf{RANK}(n,r,\tau)) = n^2 - (n-r)^2$.

Consider the following natural map Φ:

$$\mathbb{A}^{n^2-(n-r)^2} \times \mathbb{A}^k \supset \mathsf{RANK}(n,r,\tau) \times S(\pi) \xrightarrow{\Phi} M_n \cong \mathbb{A}^{n^2}, \tag{3}$$

taking (X,T_π) to $X + T_\pi$. The image of Φ is exactly $\mathsf{RIG}(n,r,\pi,\tau)$.

Also, note that $\dim(S(\pi)) = |\pi|$. We note that if there is a surjective morphism from an affine variety X to another affine variety Y, then $\dim Y \leqslant \dim X$ (we defer a formal statement to the full version [10]). Thus for $k \leqslant (n-r)^2 - 1$, we get

$$\dim(\overline{Im(\Phi)}) = \dim(\overline{\mathsf{RIG}(n,r,\pi,\tau)}) \leqslant n^2 - (n-r)^2 + k < n^2.$$

Note that

$$\mathcal{W} = \bigcup_{\tau, \pi} \overline{\mathsf{RIG}(n, r, \pi, \tau)}$$

and that completes the proof of the theorem. ∎

Thus we have proved that the set of matrices of rigidity strictly smaller than $(n-r)^2$ is contained in a proper closed affine variety of $\mathbb{A}^{n^2}$, and thus is of dimension strictly less than n^2. In other words, a *generic matrix*, i.e. a matrix that lies outside a certain proper closed affine subvariety of $\mathbb{A}^{n^2}$, is *maximally rigid*. Therefore, over an infinite field F (for instance, an algebraically closed field), there always exist maximally rigid matrices.

We now refine Valiant's argument and prove the following exact bound on the dimension of $\mathcal{W}$. The main point of the proof is a *lower bound* on $\dim(\mathcal{W})$.

THEOREM 6. *Let* $0 \leqslant r \leqslant n$ *and* $0 \leqslant k \leqslant (n-r)^2$. *Then*

$$\dim(\mathcal{W}) = n^2 - (n-r)^2 + k.$$

PROOF.   With the notation of the previous proof, we have the map

$$\Phi : \mathsf{RANK}(n, r, \tau) \times S(\pi) \to M_n.$$

defined above. Let $\mathsf{RANK}(n, \leqslant r), \mathsf{RANK}(n, r)$ be the set of $n \times n$ matrices of rank at most r and exactly r respectively. Let $S(k)$ be the set of matrices of support at most k.

Now note that $GL(n) \times GL(n)$ acts on $\mathsf{RANK}(n, \leqslant r)$ by multiplication on the left and the right, and that the action is transitive on the set of matrices with rank exactly r, which forms a Zariski open subset of $\mathsf{RANK}(n, \leqslant r)$. Therefore $\mathsf{RANK}(n, \leqslant r)$ is an irreducible algebraic variety. It is not difficult to see (for instance, from the computation below of the tangent space) that its singular locus is exactly $\mathsf{RANK}(n, \leqslant r-1)$, the set of matrices with rank less than r.

On the other hand, $S(k)$ splits into components $S(\pi)$ depending on the pattern π and is thus a union of various affine subspaces (each associated to a π of size at most k). The nonsingular elements of $S(k)$ are those which have support of size exactly k.

We can put together the maps Φ arising from various choices of τ and π to write the map

$$\widetilde{\Phi} : \mathsf{RANK}(n, \leqslant r) \times S(k) \to \mathsf{RIG}(n, r, \leqslant k).$$

We can easily see that $\widetilde{\Phi}$ is a surjective morphism of affine varieties. If we can find a nonsingular point of $\mathsf{RANK}(n, \leqslant r) \times S(k)$ for which the map on tangent spaces is injective, then the dimension of the target space $\mathsf{RIG}(n, r, \leqslant k)$ will equal $\dim \mathsf{RANK}(n, \leqslant r) + \dim S(k) = n^2 - (n-r)^2 + k$, proving the theorem. Since the map on tangent spaces is simply addition of matrices, we need to show that the subspaces do not intersect non-trivially and that would complete the proof of the theorem. For any smooth point $x \in \mathsf{RANK}(n, r)$, the smooth locus of $\mathsf{RANK}(n, \leqslant r)$, we will find a pattern π of size k and $y \in S(\pi)$ for which the tangent spaces at x and y intersect transversely.

Assume first that the point x is $\begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}$. We choose the pattern π to lie completely in the bottom right hand block of size $(n-r) \times (n-r)$, and choose any smooth point y of $S(\pi)$ (i.e. having all k entries nonzero).

The tangent space of x is $\begin{pmatrix} * & * \\ * & 0 \end{pmatrix}$.

That is, it consists of the subspace of M_n consisting of matrices with arbitrary entries except in the lower $(n-r) \times (n-r)$ block, which is constrained to be the zero submatrix. The dimension of the tangent space is $r^2 + 2r(n-r) = n^2 - (n-r)^2$, as expected. The tangent space of y is $\begin{pmatrix} 0 & 0 \\ 0 & *_\pi \end{pmatrix}$ where $*_\pi$ means that the entries in positions of π are arbitrary, and the other entries are zero.

It's clear that these two tangent spaces intersect transversely.

Now, we need to show this for a more general $x \in \text{RANK}(n,r)$. Assume that the top left $r \times r$ minor of x is nonsingular (else we can multiply by permutation matrices on left and right, noting that permutations just shuffle the various $S(\pi)$ for $|\pi| = k$).

The first r columns of x are independent and span the column space of x, so there exists a matrix $g = \begin{pmatrix} I_r & * \\ 0 & I_{n-r} \end{pmatrix}$ such that xg has the form $\begin{pmatrix} * & 0 \\ * & 0 \end{pmatrix}$. Then using that the first r rows of xg are independent and span its row space, we can find an invertible matrix $h = \begin{pmatrix} * & 0 \\ * & I_{n-r} \end{pmatrix}$ such that $hxg = \begin{pmatrix} I_r & 0 \\ 0 & 0 \end{pmatrix}$. The tangent space of x is $h^{-1}\begin{pmatrix} * & * \\ * & 0 \end{pmatrix}g^{-1}$. We need to show this does not intersect $S(\pi)$ for some π. That is, $\begin{pmatrix} * & * \\ * & 0 \end{pmatrix}$ does not intersect $h\begin{pmatrix} 0 & 0 \\ 0 & *_\pi \end{pmatrix}g$ except in zero. But this follows from the fact that the latter is a matrix of the same form (in fact, multiplication by h and g leave any element of $S(\pi)$ unchanged). $\blacksquare$

Remarks: *A similar argument or line of study - though in the projective setting - is also found in [11]. Our formalism and proofs seem clearer and simpler. Our theorem is also very explicit.*

2.3   Rigid Matrices over the field of Complex Numbers

Recall that to say that the rigidity of a matrix A for target rank r is at least k, it suffices to prove that the matrix A is not in $\mathcal{W}(n,r,\leqslant (k-1))$. We use this idea to achieve the maximum possible lower bound for the rigidity of a family of matrices over the field of complex numbers $\mathbb{C}$. As a matter of fact, we obtain matrices with real algebraic entries with rigidity $(n-r)^2$.

THEOREM 7. *Let $\delta(n) = n^{4n^4}$. Let $p_{i,j}$ for $1 \leqslant i, j \leqslant n$ be distinct primes such that $p_{i,j} > \delta(n)$. Let $K = \mathbb{Q}(\zeta_{1,1}, \ldots, \zeta_{n,n})$ where $\zeta_{i,j} = e^{2\pi i / p_{i,j}}$. Let $A(n) := (\zeta_{i,j}) \in M(n,K)$. Then, for any field L containing K,*

$$\text{Rig}(A(n), r, L) = (n-r)^2.$$

PROOF.   For simplicity, we will index the $\zeta_{i,j}$ by ζ_α for $\alpha = 1$ to n^2, and similarly p_α. We prove the theorem by showing that

$$A(n) \notin \mathcal{W}(n, r, \leqslant (n-r)^2 - 1)(L).$$

Thus it is sufficient to prove that

$$A(n) \notin \mathcal{W}(n, r, \pi)(L)$$

for any pattern π with $|\pi| = (n-r)^2 - 1$. Let π be any such pattern. To simplify notation, let us define, $\mathcal{W} := \mathcal{W}(n, r, \pi)(L)$. By Theorem 4 we have:

$$\dim(\mathcal{W}) \leqslant \dim(\mathcal{W}(n, r, \leqslant (n-r)^2 - 1)) \leqslant (n^2 - 1) < n^2.$$

Equivalently (by Hilbert's Nullstellensatz),

$$EI(n, r, \pi) \neq (0).$$

Proving that $A(n) \notin \mathcal{W}$ is equivalent to showing the existence of a $g \in EI(n, r, \pi)$ such that $g(A(n)) \neq 0$. We produce such a g using the following theorem:

THEOREM 8.([1], page 6, Theorem 4) *Let $I = \langle f_1, \ldots, f_s \rangle$ be an ideal in the polynomial ring $F[Y]$ over an infinite field F, where $Y = \{y_1, \ldots, y_m\}$. Let d be the maximum total degree of the generators f_i. Let $Z = \{y_{i_1}, \ldots, y_{i_\ell}\} \subseteq Y$ be a subset of indeterminates of Y. If $I \cap F[Z] \neq (0)$ then there exists a non-zero polynomial $g \in I \cap F[Z]$ such that, $g = \sum_{i=1}^{s} g_i f_i$, with $g_i \in F[Y]$ and $\deg(g_i f_i) \leqslant (\mu + 1)(m + 2)(d^\mu + 1)^{\mu+2}$, where $\mu = min\{s, m\}$.*

Let us apply Theorem 8 to our case - in the notation of this theorem our data is as follows: $F := \mathbb{Q}, Y := \{x_1, \ldots, x_{n^2}, t_1, \ldots, t_k\}, Z := \{x_1, \ldots, x_{n^2}\}, \Sigma_{r+1} :=$ set of all minors of size $(r+1)$, $f_\tau := \det((X + T_\pi)_\tau)$ for $\tau \in \Sigma_{r+1}$, here by Y_τ we denote the τ-th minor of Y, and $I := I(n, r, \pi) = \langle f_\tau : \tau \in \Sigma_{r+1} \rangle$ as defined in (1).

Furthermore, we have:

$$
\begin{aligned}
m &= n^2 + (n-r)^2 - 1 \leqslant 2n^2 - 2 \\
\mu &= \min \left\{ n^2 + (n-r)^2 - 1, \binom{n}{r+1}^2 \right\} \\
&\leqslant n^2 + (n-r)^2 - 1 \leqslant 2n^2 - 2, \\
d &= r + 1 \leqslant n, \\
I \cap F[Z] &= EI(n, r, \pi) \neq (0).
\end{aligned}
$$

By Theorem 8 there exists a

$$g \neq 0 \in EI(n, r, \pi) \subseteq \mathbb{Q}[x_1, \ldots, x_{n^2}]$$

such that

$$\deg(g) \leqslant (2n^2 - 1)(2n^2)(n^{2n^2-2} + 1)^{2n^2} < n^{4n^4} = \delta(n).$$

We will now apply the following Lemma 9, which we prove later, to this situation.

LEMMA 9. *Let N be a positive integer. Let $\theta_1, \cdots, \theta_m$ be m algebraic numbers such that for any $1 \leqslant i \leqslant m$, the field $\mathbb{Q}(\theta_i)$ is Galois over $\mathbb{Q}$ and such that*

$$[\mathbb{Q}(\theta_i) : \mathbb{Q}] \geqslant N \quad \text{and}$$

$$\mathbb{Q}(\theta_i) \cap \mathbb{Q}(\theta_1, \ldots, \theta_{i-1}, \theta_{i+1}, \ldots, \theta_m) = \mathbb{Q}.$$

Let $g(\underline{x}) \neq 0 \in \mathbb{Q}[x_1, \ldots, x_m]$ such that $\deg(g) < N$. Then,

$$g(\theta_1, \ldots, \theta_m) \neq 0.$$

Let us set $m = n^2, N = \delta(n), l := \deg(g) \leqslant N$ in Lemma 9. It is now easy to check that

$$[\mathbb{Q}(\zeta_\alpha) : \mathbb{Q}] = p_\alpha - 1 \geqslant \delta(n) = N$$

and

$$\mathbb{Q}(\zeta_\alpha) \cap \mathbb{Q}(\zeta_1, \ldots, \zeta_{\alpha-1}, \zeta_{\alpha+1}, \ldots, \zeta_{n^2}) = \mathbb{Q}.$$

The latter follows from the fact that the prime p_α is totally ramified in $\mathbb{Q}(\zeta_\alpha)$ and is unramified in $\mathbb{Q}(\zeta_1, \ldots, \zeta_{\alpha-1}, \zeta_{\alpha+1}, \ldots, \zeta_{n^2})$; see Theorem 4.10 in [17]. Thus Lemma 9 is applicable and we get:

$$g(\zeta_1, \ldots, \zeta_{n^2}) \neq 0.$$

To complete the argument (for Theorem 7), now we prove Lemma 9.

Proof of Lemma 9 : By induction on m. For $m = 1$ this is trivial. Now suppose that the statement is true when the number of variables is strictly less than m. Assuming that the statement is not true for m, we will arrive at a contradiction. This will prove the Lemma.

Let $g \in \mathbb{Q}[\underline{x}]$ with $l := \deg(g) < N$ be such that

$$g(\theta_1, \ldots, \theta_m) = 0,$$

with θ_i, $1 \leqslant i \leqslant m$, satisfying the conditions as in the theorem. Since the statement is true for any $(m-1)$ number of variables, without loss of generality, we can assume that all the variables and hence x_m appears in g. Let us denote x_m by x. Let us write

$$g(x_1, \ldots, x_m) = \sum_{i=0}^{l} f_i(x_1, \ldots, x_{m-1}) x^{l-i}.$$

Note that $l < N$ and $\deg(f_i) < N$ for $0 \leqslant i \leqslant l$. Since $g \neq 0$, for some i, $0 \leqslant i \leqslant l$ the polynomial $f_i \neq 0$. Thus, by the inductive hypothesis,

$$f_i(\theta_1, \ldots, \theta_{m-1}) \neq 0.$$

Thus $g(\theta_1, \ldots, \theta_{m-1})(x) \neq 0 \in \mathbb{Q}(\theta_1, \ldots, \theta_{m-1})[x]$. This implies that θ_m satisfies a non-zero polynomial over $\mathbb{Q}(\theta_1, \ldots, \theta_{m-1})$ of degree $\leqslant l < N$. Thus:

$$[\mathbb{Q}(\theta_1, \ldots, \theta_m) : \mathbb{Q}(\theta_1, \ldots, \theta_{m-1})] \leqslant l < N. \tag{4}$$

On the other hand, since $\mathbb{Q}(\theta_m) \cap \mathbb{Q}(\theta_1, \ldots, \theta_{m-1}) = \mathbb{Q}$ and the fields $\mathbb{Q}(\theta_i)$ are Galois over $\mathbb{Q}$, it can be concluded by the property of such extensions ([12] Theorem 1.12, page 266) that

$$[\mathbb{Q}(\theta_1, \ldots, \theta_{m-1})(\theta_m) : \mathbb{Q}(\theta_1, \ldots, \theta_{m-1})] = [\mathbb{Q}(\theta_m) : \mathbb{Q}] \geqslant N.$$

This contradicts (4) above and that proves the lemma.

This concludes the proof of Theorem 7.■

Note that Theorem 7 is true for any family of matrices $A(n) = [\theta_{i,j}]$ provided the $\theta_{i,j}$ satisfy Lemma 9. Hence, we have

COROLLARY 10. *Let $A(n) := (\zeta_{i,j} + \overline{\zeta_{i,j}})$, where $\zeta_{i,j}$ are primitive roots of unity of order $p_{i,j}$ such that $p_{i,j} - 1 \geqslant 2\delta(n)$ (here $\overline{\zeta_{i,j}}$ denotes the complex conjugate of $\zeta_{i,j}$). Then, $A(n) \in M(n, \mathbb{R})$ has $\mathrm{Rig}(A(n), r) = (n - r)^2$.*

3   Reduction to Determinantal Ideals

In this section, we show that the natural decomposition of the rigidity varieties $\mathcal{W}(n, r, \leqslant k) = \cup_{|\pi|=k} \mathcal{W}(n, r, \pi)$ is indeed a decomposition into *irreducible* affine algebraic varieties. In fact, these components turn out to be varieties defined by elimination ideals of determinantal ideals generated by all the $(r + 1) \times (r + 1)$ minors.

To show the decomposition, we will continue to use the notation from Section 2. Consider the matrix $X + T_\pi$. Let $x = \{x_1, \ldots, x_{n^2}\} = x_{\bar\pi} \cup x_\pi$, where x_π is the set of variables that are indexed by π and $x_{\bar\pi}$ is the set of remaining variables.

Let

$$J := I(n, r, \pi) = \left\langle Minors_{(r+1) \times (r+1)}(X + T_\pi) \right\rangle$$

be the ideal of $\mathbb{Q}[x, t] = \mathbb{Q}[x_\pi, x_{\bar\pi}, t_\pi]$ generated by the $(r + 1) \times (r + 1)$ minors of $X + T_\pi$. Let

$$
\begin{aligned}
J_1 &:= J \cap \mathbb{Q}[x_\pi, x_{\bar\pi}] \subseteq \mathbb{Q}[x_1, \ldots, x_{n^2}], \\
J_2 &:= J_1 \cap \mathbb{Q}[x_{\bar\pi}], \\
I_{r+1} &:= \left\langle Minors_{(r+1) \times (r+1)}(X) \right\rangle \subseteq \mathbb{Q}[x], \quad \text{and} \\
EI_{r+1} &:= I_{r+1} \cap \mathbb{Q}[x_{\bar\pi}] \subseteq \mathbb{Q}[x_{\bar\pi}].
\end{aligned}
$$

Notice that since J_1 is the elimination ideal of J w.r.t. eliminating variables t_π, a matrix A lies in $\mathcal{W}(n, r, \leqslant k) = \overline{\mathrm{RIG}(n, r, \leqslant k)}$ if and only if its entries lie in the variety defined by the ideal J_1. Also, I_{r+1} is the ideal generated by the $(r + 1) \times (r + 1)$ minors of X and EI_{r+1} its elimination ideal for the rational ring generated by the variables $x_{\bar\pi}$.

PROPOSITION 11. $J_1 = J_2\mathbb{Q}[x]$ *(the ideal generated by J_2 in $\mathbb{Q}[x]$) and $J_2 = EI_{r+1}$. In particular, $EI(n, r, \pi) = EI_{r+1}\mathbb{Q}[x]$ considered as ideals in $\mathbb{Q}[x]$.*

PROOF.   First, notice that in the $(r + 1) \times (r + 1)$ minors of $X + T_\pi$, the variable $t_{i,j}$, for $(i, j) \in \pi$, always occurs in combination with $x_{i,j}$ as $t_{i,j} + x_{i,j}$. Therefore, eliminating the variables t_π will also automatically eliminate the variables x_π, giving the equality of the generators of the ideals J_1 and J_2. Therefore $J_1 = J_2\mathbb{Q}[x]$. More formally, consider the isomorphism between the two coordinate rings $\phi : \mathbb{Q}[x_\pi, x_{\bar\pi}, t_\pi]$ and $\mathbb{Q}[x_\pi, x_{\bar\pi}, t_\pi]$ defined by letting $\phi(t_{i,j}) = x_{i,j} + t_{i,j}$ for each $(i, j) \in \pi$ and $\phi(x_{i,j}) = x_{i,j}$ for all $(i, j) \notin \pi$. The ideal

$J_1 = J \cap \mathbb{Q}[x_\pi, x_{\bar{\pi}}] \subseteq \mathbb{Q}[x_1, \ldots, x_{n^2}]$ must equal the ideal $\phi(\phi^{-1}(J) \cap \phi^{-1}\mathbb{Q}[x_1, \ldots, x_{n^2}])$, since ϕ is an isomorphism. But $\phi^{-1}(J)$ is generated by matrices only involving the variables of t_π and $x_{\bar{\pi}}$, whereas $\phi^{-1}\mathbb{Q}[x_1, \ldots, x_{n^2}]) = \mathbb{Q}[x_1, \ldots, x_{n^2}]$, so that $\phi^{-1}(J) \cap \phi^{-1}\mathbb{Q}[x_1, \ldots, x_{n^2}]$ is generated by polynomials only involving the variables of $x_{\bar{\pi}}$. Therefore $\phi^{-1}(J_1) = \phi^{-1}(J) \cap \phi^{-1}\mathbb{Q}[x_1, \ldots, x_{n^2}] = J_2\mathbb{Q}[x]$ and taking the image under ϕ, we get $J_1 = J_2\mathbb{Q}[x]$.

The equation $J_2 = EI_{r+1}$ follows from similar considerations, noting that the variables $x_{i,j}$ for $(i,j) \in \pi$ always occur in the combination $x_{i,j} + t_{i,j}$. Therefore eliminating them eliminates $t_{i,j}$ as well. More formally, consider the isomorphism $\psi : \mathbb{Q}[x_\pi, x_{\bar{\pi}}, t_\pi] \rightarrow \mathbb{Q}[x_\pi, x_{\bar{\pi}}, t_\pi]$ defined by letting $\psi(x_{i,j}) = x_{i,j} + t_{i,j}$ for each $(i,j) \in \pi$, while $\psi(t_{i,j}) = t_{i,j}$ for $(i,j) \in \pi$ and $\psi(x_{i,j}) = x_{i,j}$. Then again we have $J_2 = J_1 \cap \mathbb{Q}[x_{\bar{\pi}}] = J \cap \mathbb{Q}[x_{\bar{\pi}}] = \psi(\psi^{-1}(J) \cap \psi^{-1}(\mathbb{Q}[x_{\bar{\pi}}])) = \phi(I_{r+1}\mathbb{Q}[x, t_\pi] \cap \mathbb{Q}[x_{\bar{\pi}}]) = \phi(EI_{r+1}) = EI_{r+1} \subset \mathbb{Q}[x_{\bar{\pi}}]$. ∎

The following is a well-known theorem; see [3, Chapter 2].

THEOREM 12. *Let* $\mathrm{RANK}(n, \leqslant r)$ *be the set of all rank* $\leqslant r$ *matrices of* $M_n \cong \mathbb{A}^{n^2}$. *Then*
- $I(\mathrm{RANK}(n, \leqslant r)) = I_{r+1}$ *and* $\mathrm{RANK}(n, \leqslant r) = V(I_{r+1})$.
- I_{r+1} *is a prime ideal of* $\mathbb{Q}[X]$. *In particular,* $\mathrm{RANK}(n, \leqslant r)$ *is an irreducible variety.*

From Theorem 12 and Proposition 11 we get the following corollary (see [10]) for details).

COROLLARY 13. *In the natural decomposition* $\mathcal{W}(n, r, \leqslant k) = \cup_{|\pi|=k} \mathcal{W}(n, r, \pi)$, *the* $\mathcal{W}(n, r, \pi)$ *are irreducible varieties.*

4 Semicontinuity of Rigidity

Intuitively, if a function is (lower) semicontinuous at a given point, then within a small neighborhood of that point the function is nondecreasing. (See the full version [10] of the paper for a formal treatment of the material in this section). The rank function of a matrix, for example, is a lower semicontinuous function on the space of all $n \times n$ complex matrices. It is possible to construct give examples (we defer this to the full version [10]) to show that the rigidity function is not semicontinuous in general. However, it seems to have semicontinuity property at some interesting matrices. In particular, the matrices $A(n)$ from Theorem 7 have an open neighborhood around them within which the rigidity function is constant. This is a direct consequence of their very construction since they are outside the closed sets $\mathcal{W}(n, r, \leqslant (n-r)^2 - 1)$. These examples motivate us to study the properties of the Euclidean closure and Zariski closure of the set $\mathrm{RIG}(n, r, \leqslant k)(\mathbb{C})$. In fact, we are able to argue that these two coincide.

PROPOSITION 14. *The Euclidean Closure of* $\mathrm{RIG}(n, r, \leqslant k)(\mathbb{C})$ *equals its Zariski Closure.*

PROOF.   Recall that we can write $\mathrm{RIG}(n, r, \leqslant k) = \bigcup_{\pi, |\pi|=k} \mathrm{RIG}(n, r, \pi)$. Thus, to prove the proposition, it is sufficient to prove that for any pattern π, the Euclidean closure of $\mathrm{RIG}(n, r, \pi)$ equals its Zariski Closure. By Closure Theorem, there exists a subvariety V strictly contained in $\mathcal{W} := \overline{\mathrm{RIG}(n, r, \pi)}$ such that $\mathcal{W}(\mathbb{C}) - V(\mathbb{C}) \subseteq \mathrm{RIG}(n, r, \pi)(\mathbb{C}) \subseteq \mathcal{W}(\mathbb{C})$. Since $\mathcal{W}(\mathbb{C})$ is closed in the Euclidean topology, we will done if we prove that the Euclidean closure of $\mathcal{W}(\mathbb{C}) - V(\mathbb{C})$ is $\mathcal{W}(\mathbb{C})$. This is precisely the statement of the following lemma from [20], which we state below for easy reference. Also note that, by Corollary 13, W is an irreducible variety for every pattern π and hence the lemma is applicable.

LEMMA 15. ([20, Lemma 1, page 124]) *If X is an irreducible algebraic variety and Y a proper subvariety of X then the set $X(\mathbb{C}) - Y(\mathbb{C})$ is dense in $X(\mathbb{C})$.*

References

[1] A. Bernasconi, E. W. Mayr, M. Mnuk, and M. Raab. Computing the Dimension of a Polynomial Ideal. http://www14.informatik.tu-muenchen.de/personen/raab/, 2002.

[2] W. D. Brownawell. Bounds on the degrees of Nullstellensatz. *Annals of Mathematics*, 126:577–592, 1987.

[3] W. Bruns and U. Vetter. *Determinantal Rings*, volume 1327 of *Lect. Notes in Math.* 1980.

[4] M. Cheraghchi. On Matrix Rigidity and the Complexity of Linear Forms. *Electronic Colloquium on Computational Complexity (ECCC)*, (070), 2005.

[5] B. Codenotti. Matrix Rigidity. *Linear Algebra and its Applications*, 304(1–3):181–192, 2000.

[6] D. Cox, J. Little, and D. O'Shea. *Ideals, Varieties,and Algorithms*. Under Graduate Text-books in Mathematics. 3rd edition, 2007.

[7] J. Forster. A Linear Lower Bound on the Unbounded Error Probabilistic Communica-tion Complexity. *Journal of Computer and System Sciences*, 65(4):612–625, 2002.

[8] J. Friedman. A Note on Matrix Rigidity. *Combinatorica*, 13(2):235 – 239, 1993.

[9] J. Kollar. Sharp Effective Nullstellensatz. *Jl. of American Math. Soc.*, 1(4):963–975, 1988.

[10] A. Kumar, S. V. Lokam, V. Patankar, and J. M. N. Sarma. Using Elimination Theory to Construct Rigid Matrices. Manuscript, 2009.

[11] J. M. Landsberg, J. Taylor, and N. K. Vishnoi. The Geometry of Matrix Rigidity. Tech-nical Report GIT-CC-03-54, Georgia Institute of Technology, 2003.

[12] S. Lang. *Algebra*. Springer-Verlag, revised third edition, 2004.

[13] N. Linial and A. Shraibman. Learning complexity vs communication complexity. *Com-binatorics, Probability and Computing*, 18(1-2):227–245, 2009.

[14] S. V. Lokam. On the Rigidity of Vandermonde matrices. *Theoretical Computer Science*, 237:477–483, 2000.

[15] S. V. Lokam. Spectral Methods for Matrix Rigidity with Applications to Size-Depth Tradeoffs and Communication Complexity. *Jl. of Comp. Syst. Sci.*, 63(3):449–473, 2001.

[16] S. V. Lokam. Quadratic Lowerbounds on Matrix Rigidity. In *Proc. of Internatial Conf. on Theory and Applications of Models of Computation*, volume 3959 of *LNCS*, 2006.

[17] W. Narkiewicz. *Elementary and Analytic Theory of Algebraic Numbers*, volume XI of *Springer Monographs in Mathematics*. Springer, 2004.

[18] R. Paturi and P. Pudlák. Circuit Lower Bounds and Linear Codes. *Teoria slozhnosti vychislenij IX*, 316:188–204, 2004. ECCC Techreport : TR04-04.

[19] A. A. Razborov. On Rigid Matrices. Manuscript, (Russian), 1989.

[20] I. R. Shafarevich. *Varieities in Projective Space*, volume 1 of *Basic Algebraic Geometry*. Springer Verlag, second edition, 1994.

[21] D. A. Spielman, V. Stemann, and M. A. Shokhrollahi. A Remark on Matrix Rigidity. *Information Processing Letters*, 64(6):283 – 285, 1997.

[22] L. G. Valiant. Graph Theoretic Arguments in Low Level Complexity. volume 53 of *Lecture Notes in Computer Science*, pages 162–176. Springer Verlag, 1977.