Short polynomials and where to find them

Thomas Kahle

June 15, 2022

1 Introduction

We consider ideals in a polynomial ring $k[x_1, \dots, x_n]$, where k is some field. The two main cases we have in mind are $k = \mathbb{Q}$ and $k = \mathbb{C}$ because in $\mathbb{Q}$ we can compute with exact arithmetic and because $\mathbb{C}$ is algebraically closed. The questions are interesting for any field, though. We use some standard notation such as *monomial notation* $x^u := \prod_{i=1}^n x_i^{u_i}$ with $u \in \mathbb{Z}^n$. As visible with the negative exponents, in some cases we might also switch to a Laurent polynomial ring $k[x_1^{\pm}, \dots, x_n^{\pm}]$

Definition 1.1. A polynomial is *t-short* if it has at most t-terms. A *monomial* is a 1-short polynomial. A *binomial* is a 2-short polynomial.

Remarks.

- We always work in the monomial basis.
- If there is no t, then what „short“ means can depend a lot on the context.
- Other names that have appeared in the literature are *fewnomials* [Kho91] (although that book goes in a very different direction), *segmentonomials* [BG02], ...

The following is an open problem to be approached in this lecture. We present a solution for the $t = 2$ case, i.e. the search for binomials.

Problem 1.2. *Let $I \subset \mathbb{Q}[x_1, \dots, x_n]$ be an ideal and $t > 2$ an integer. Is there an algorithm that decides if I contains a t-short polynomial?*

Remarks.

- We are not yet searching for an efficient algorithm. *Any* algorithm will do. Feel free to compute as many Gröbner bases on the way as necessary.

If the decision problem can be solved, it would be interesting to consider auxiliary problems such as the optimization problem "What is the shortest length of a polynomial in I?" or "What are all the shortest polynomials in I?"

The $t = 1$ question, whether I contains a monomial can be readily solved using Gröbner bases.

Proposition 1.3.

$$I \text{ contains a monomial} \quad \Leftrightarrow \quad I : (\prod_{i=1}^{n} x_i)^{\infty} = k[x_1, \ldots, x_n]$$

Here $I : f = \{ g \in k[x_1, \ldots, x_n] : fg \in I \}$ is the ideal quotient and the exponent ∞ means to repeat this operation to stabilization. All the operations used in this example are computationally solved by means of Gröbner bases. The proof is an easy exercise.

Example. Consider the binomial ideal $I = \langle x^2 - xy, xy + y^2 \rangle \in \mathbb{Q}[x, y]$. These two generators conflict in degree three. Multiplying the first generator by y yields $x^2y = xy^2$ while multiplying the second generator by x yields $x^2y = -y^2x$ (these are equations in the quotient ring). Both can only be true at the same time, if $xy^2 = x^2y = 0$. In fact all degree 3 monomials are contained in I.

There is also a geometric viewpoint which is often useful to think about.

Problem 1.4. *Pick your favorite algebraic variety (or just favorite set) X. What is the shortest polynomial that vanishes on X?*

This is asking for the shortest polynomial in the ideal $I(X)$ of all polynomials vanishing on X. Such a vanishing ideal $I(X)$ is always a radical ideal ($f^s \in I \Rightarrow f \in I$) and in computational algebra it is often advantageous to work with radical ideals. If one tries to prove that no short polynomial exists in I, then one can also try to prove it geometrically and thus for $\sqrt{I}$. Since $I \subset \sqrt{I}$ this suffices and it might be easier because radical ideals are easier. On the other hand, being larger, $\sqrt{I}$ might contain shorter polynomials.

Here is our plan for the lectures.

Contents

2 Matroid theory and linear algebra for bounded degrees

Our first approach is based on the observation that linear algebra provides algorithms and methods to find a shortest vector in a linear space. This is elimination or row reduction, but simple computation of row echolon forms is not enough. In any case, an ideal in the polynomial k-algebra is a linear space and googling for *sparse vector in linear space* will yield many results. These go in a somewhat different direction. Sparsity in linear algebra is the main way to deal with large matrices, i.e. matrices that are so large that one cannot store them as an array in the memory. One can still work with these, if there are enough zeros and only non-zero entries and their positions are stored. This is sparse linear algebra. Algorithmic questions in this area are usually not about the existence of algorithms, but about their complexity. See [McC83] for an early survey.[1] Quite generally, this area of linear algebra is very much connected to mathematical optimization.

We will start from scratch. Let $L \subseteq k^n$ be some linear space with $\dim(L) = k$ and $v_1, \dots, v_k$ a basis. Using orthogonal direct sums (i.e. the computation of kernels) one can find a matrix $A \in k^{m \times n}$ so that $\ker(A) = L$, $m = n - k$ is the minimum number of rows necesary. The columns of A define a vector matroid M which is independent of the choices and depends only on the space L. Finding short vectors is finding the *circuits*[2] of this matroid, which are the inclusion minimimal elements of $\{ \operatorname{supp}(v) : v \in L \}$. Software that can compute the circuits of a matrix includes TOPCOM, Polymake and sage.

With this in mind, Problem 1.2 can be solved as soon as there is a computable degree bound[3] for the t-short polynomial. With such a degree bound, finding short polynomials is finding short vectors in a finite-dimensional vector space and thus we get the following proposition.

Proposition 2.1. *There is an algorithm which given an ideal I and integers d, t decides if I contains a t-short polynomial of degree at most d.*

[1] Chapter 3 is titled „Making matrices sparser“!

[2] The circuits are one of the many cryptomorphic ways to define a matroid. Precisely they are a set $\mathcal{C}$ of subsets $C \subset [n]$ such that
(i) $\emptyset \notin \mathcal{C}$
(ii) if $C_1, C_2 \in \mathcal{C}$ and $C_1 \subseteq C_2$, then $C_1 = C_2$
(iii) if $C_1 \neq C_2 \in \mathcal{C}$ and $e \in C_1 \cap C_2$, then there is a $C_3 \in \mathcal{C}$ such that $C_3 \subseteq (C_1 \cup C_2) \setminus \{ e \}$.
These axioms are satisfied for the supports of shortest vectors in any linear space. Here (iii) is the eliminiation property in the kernel and (ii) is the desired minimality. Circuits are called circuits because they are also the shortest cycles in matroids coming from graphs.

[3] This means a theorem of the sort „If I contains a t-short polynomial then it contains a t-short polynomial of degree at most d“ where d is computable from I.

Proof. Let $I_{\leq d} = I \cap k[x_1, \ldots, x_n]_{\leq d}$ consist of all polynomials in I of degree at most d. Compute a basis of the finite-dimenisonal vector space $I_{\leq d}$. From this basis, compute the basis of an orthogonal direct complement of $I_{\leq d}$ (i.e. the right kernel of the matrix). Write the basis of the complement as rows of coefficients (with respect to the monomial basis) of a matrix A. Compute the finitely many circuits of A using e.g. [Min76]. Decide from the finite list, if there exists a circuit of length $\leq t$. □

Example. We consider a univariate example (composed by Niclas Niederdrenk). Let $f = (x+1)(x^2+x+1) = x^3+2x^2+2x+1$ and $I = \langle f \rangle$ be the principal ideal generated by f. Since f is a product of the 2nd and 3rd cyclotomic polynomials and $2, 3$ are divisors of 6, we get that f divides $x^6 - 1$ which consequently lies in I. We verify this in sage with the following code

```
A5 = matrix (QQ,[[1,2,2,1,0,0],[0,1,2,2,1,0],
                 [0,0,1,2,2,1]])
A6 = matrix (QQ,[[1,2,2,1,0,0,0],[0,1,2,2,1,0,0],
                 [0,0,1,2,2,1,0],[0,0,0,1,2,2,1]]
S5 = A5.right_kernel(); S5
S6 = A6.right_kernel(); S6
M5 = Matroid (S5.basis_matrix())
for c in M5.circuits(): print(c)
M6 = Matroid(S6.basis_matrix())
for c in M6.circuits(): print(c)
```

The last line of output shows a circuit $\{0, 6\}$ which corresoponds to the binomial. That all circuits of `S5` have four elements shows that this is the lowest degree binomial in I.

Remark. Duality is very important in matroid theory. The algorithm in Proposition 2.1 can also work directly with the basis of $I_{\leq d}$. In this case one computes the hyperplanes (or maximal flats) the vector matroid corresponding to the basis. They hyperplanes are the complements of the cocircuits, which means that one tries to find large hyperplanes whose complements show the supports of short polynomials. We continue the sage session from Example 2 as follows:

```
M6d = Matroid (A6)
for h in M6d.hyperplanes(): print(h)
print ("-----")
for c in M6d.cocircuits(): print(c)
```

So, after all, we only need a computable degree bound for a shortest polynomial to solve Problem 1.2. But this might be harder to get than one thinks. Let's consider two examples:

Example. Let F_n be the n-th Fibonacci number and

$$I_n = \langle y - F_n x - F_{n-1}, x^2 - x - 1 \rangle.$$

This ideal is generated by a quadric and a linear trinomial and it contains the binomial $x^n - y$ which can be shown by the reduction

$$x^n \to x^{n-2}(x+1) = x^{n-1} + x^{n-2} \to x^{n-2} + 2x^{n-3} + x^{n-4} \to \cdots \to y.$$

This is the lowest degree binomial in I. This means that

- The computable degree bound, if one exists, needs to depend on the *coefficients* in the generators.
- No classical measures of complexity of an ideal such as generating degree, Castelnuovo–Mumford regularity or anything derived from primary decomposition can yield a degree bound.
- Homogenization plays no role as the example works also after homogenization and keeps its properties.

It meets the eye that the coefficients in the example are large, but this need not be the case as an improved example shows:

Example. For any $n \in \mathbb{N}$, let $I = \langle (x-z)^2, nx - y - (n-1)z) \rangle \subset \mathbb{Q}[x,y,z]$. The Castelnuovo–Mumford regularity of I is 2 and it is primary over $\langle x - z, y - z \rangle$. The binomial $x^n - yz^{n-1}$ is contained in I because an elementary computation shows that

$$x^n - yz^{n-1} = \sum_{k=0}^{n-2} (n-k-1)x^k z^{n-k-2}(x-z)^2 + z^{n-1}(nx - y - (n-1)z) \in I.$$

There is no binomial of degree less than n in I. To see this, consider the differential operators $D_1 = \partial_x + n\partial_y$ and $D_2 = (1-n)\partial_y + \partial_z$. Any element $f \in I$ satisfies $f(1,1,1) = 0$, $(D_1 f)(1,1,1) = 0$ and $(D_2 f)(1,1,1) = 0$ as both generators have this property. Assume that I contains the binomial $f = x^u - \lambda y^v$. First, note that $f(1,1,1) = 0$ implies that $\lambda = 1$. Further, $(D_1 f)(1,1,1) = 0$ and $(D_2 f)(1,1,1) = 0$ give two linear conditions on the vector $u - v$, which imply that $u - v = m(n, -1, 1-n)$ for some $m \in \mathbb{Z}$. By exchanging u and v we may assume that $m > 0$, so it follows that $f = x^{mn} - y^m z^{m(n-1)}$. In particular, there is no binomial of degree less than n in I.

Remark. Even if no degree bound is known, one can consider the family of matroids M_d arising from vector spaces $I_{\leq d}$. By a result of MacLagan and Rincón this data structure is equivalent to the *tropical scheme* defined by a homogeneous ideal I [MR19].

3 Finding binomials in polynomial ideals

In this section we consider the question, to compute all binomials in an ideal $I \subset \mathbb{Q}[x_1, \dots, x_n]$ and in particular, to decide if there are some. Everything is based on [JKK17].

Remarks.

- The question if I is generated by binomials is much easier to solve (only one Gröbner basis by [ES96]).
- It is also easy to decide if I contains one specific binomial $x^u - \lambda x^v$ given u, v but unknown λ. To do this, compute normal forms of x^u and x^v with Gröbner bases and decide if they are $\mathbb{Q}$-multiples of each other.
- The univariate case was solved by [GRT10]. Let $f \in \mathbb{Q}[x]$ be a non-constant monic polynomial with constant term 1 (a situation that can be achieved with tricks). Then $\langle f \rangle$ contains a binomial if and only if f is a squarefree product of cyclotomic polynomials.
- By Proposition 1.3, we can always assume that I contains no monomials. More generally, we just pass to the Laurent ring. Define $T := \mathbb{Q}[x_1^{\pm}, \dots, x_n^{\pm}]$. For any $I \in \mathbb{Q}[x_1, \dots, x_n]$ it holds that

$$IT \cap \mathbb{Q}[x_1, \dots, x_n] = I : (\prod_{i=1}^{n} x_i)^{\infty}.$$

Our line of attack is to use tropical geometry to reduce to the case where I is an Artinian ideal and that case was basically already solved by number theorists. Let's start to define what we are after.

Definition 3.1. Let $I \subset S$ be an ideal. The *binomial part* $\mathrm{Bin}(I)$ *of* I is the k-subspace of I spanned by all binomials in I.

Proposition 3.2. *(i) The binomial part of any ideal is a binomial ideal.*
(ii) $\mathrm{Bin}(IT) = \mathrm{Bin}(I)T$ *for any* $I \subset \mathbb{Q}[x_1, \dots, x_n]$.
(iii) If $I : (x_1 \cdots x_n) = I$, *then* $\mathrm{Bin}(I) = \mathrm{Bin}(Ik[x_1^{\pm}, \dots, x_n^{\pm}]) \cap k[x_1, \dots, x_n]$.
(iv) If $K/\mathbb{Q}$ *is any field extension, then* $\mathrm{Bin}(IK[x_1, \dots, x_n]) = \mathrm{Bin}(I)K[x_1, \cdots, x_n]$.

See [JKK17] for proofs. Hints: For (i), if $b \in \mathrm{Bin}(I)$, then b is a linear combination of binomials in I. Multiplication with an arbitrary f yields a linear combination of monomial multiples of those binomials which are also all in I. For (ii) use that any binomial in IT can be multiplied by a monomial to become a binomial in I.

Remarks.

- As a vector space, a binomial ideal is spanned by the binomials it contains. In particular, a binomial ideal equals its binomial part.

- When passing to the Laurent ring, many new binomials are created, but (ii) says that this would not generate binomials if none where there to start with. The new binomials only arise from moving existing binomials around. In particular, I contains binomials if and only if IT does.
- I contains a binomial if and only if the extensions to the Laurent ring contains one. We can thus also extend the field for free.

We will use some *tropical geometry* which is the piece-wise linear skeleton of algebraic geometry.

Definition 3.3. Let $I \subset T$. The *tropical variety of* I is the following subcomplex of the Gröbner fan

$$T(I) = \{ w \in \mathbb{Q}^n : \text{in}_w(I) \neq T \}.$$

Here $\text{in}_w(I)$ is the lead term ideal with respect to the term order induced by $w \in \mathbb{Q}^n$, under which exponents u, v are compared as $u \leq v \Leftrightarrow \langle u, w \rangle \leq \langle v, w \rangle$.

Remarks.

- Tropical geometry encodes polyhedral aspects of algebraic geometry.
- The Gröbner fan is the polyhedral subdivision of $\mathbb{Q}^n$ which arises when decorating each point $w \in \mathbb{Q}^n$ with its corresponding initial ideal.
- Most initial ideals are monomial ideals and these are exactly those that are excluded in the tropical variety.

The crucial facts are the following:

- If b is a binomial, then $T(\langle b \rangle)$ is an ordinary hyperplane in $\mathbb{Q}^n$, perpendicular to the Newton polytope of b (or empty).
- If some binomial $b \in I$, then $T(I) \subseteq T(\langle b \rangle)$ as we are used to in algebraic geometry.

According to these, our plan of attack is to decide if the tropical variety is contained in a hyperplane, certainly a questions that can be decided algorithmically. However, this does not answer everything. It just yields the Artinian reduction.

Example. Let $I = \langle x - y \rangle \subset \mathbb{Q}[x^{\pm}, y^{\pm}]$. Our goal is to reduce the computation of $\text{Bin}(I)$ to Artinian case. The tropical variety of I is

$$T(I) = \left\{ \lambda \cdot \binom{1}{1}, \lambda \in \mathbb{Q} \right\}.$$

This means that we can consider $T(I)^{\perp} \cap \mathbb{Z}^2$ and write

$$\langle x - y \rangle = \langle xy^{-1} - 1 \rangle$$

Then $I' = \langle xy^{-1} - 1 \rangle \subset \mathbb{Q}[(xy^{-1})]$ is Artinian and the binomials in I and I' are computable from each other.

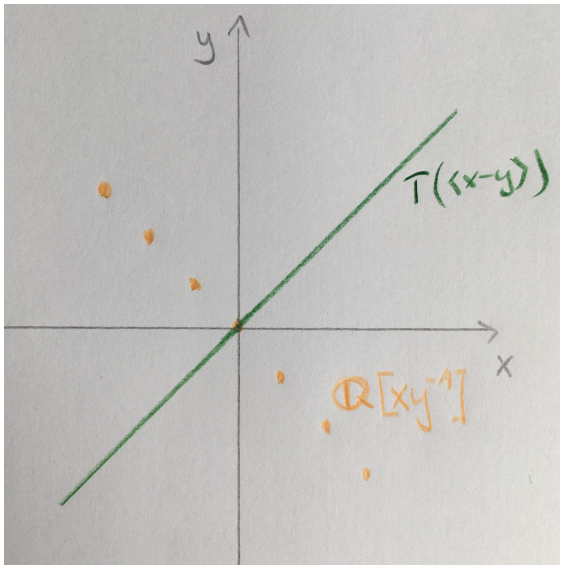

Figure 1: The situation in Example 3

Theorem 3.4. *Let k be any field, $I \subset k[x_1^{\pm}, \dots, x_n^{\pm}]$ an ideal. Then*

(i) $\mathrm{Bin}(I) = \mathrm{Bin}(I \cap k[T(I)^{\perp} \cap \mathbb{Z}^n])k[\mathbb{Z}^n]$

(ii) $I \cap k[T(I)^{\perp} \cap \mathbb{Z}^n]$ *is an Artinian ideal in* $K[T(I)^{\perp} \cap \mathbb{Z}^n]$.

Remarks.

- The proof uses the Sturmfels–Tevelev theorem which roughly says that multiplicative coordinate changes on $k[x_1^{\pm}, \dots, x_n^{\pm}]$ correspond to linear coordinate changes on the tropical variety.
- A theorem of Bieri–Groves says that if the tropical variety is zero-dimensional, then the variety is zero-dimensional.
- There is a corollary that was our first aim. If I is geometrically prime (i.e. the extension to the algebraic closure is prime), then I contains a binomial if and only if $T(I)$ is contained in a hyperplane.
  (i) $I' = I \cap k[T(I)^{\perp} \cap \mathbb{Z}^n]$ is proper Artinian.
  (ii) Identify $k[T(I)^{\perp} \cap \mathbb{Z}^n] = k[s_1^{\pm}, \dots, s_r^{\pm}]$.
  (iii) I' contains a non-constant univariate Laurent polynomial in each s_i.
  (iv) Since I is prime, the binomial factors are contained in I.
  This does not work if $k = \mathbb{Q}$ is not algebraically closed. For example $I = \langle (x-1)(x-2) \rangle$ does not contain a binomial, but its tropical variety is $\{0\}$.
- Algorithms: Computing the tropical variety is possible. We give a somewhat more efficient algorithm which does not compute the whole tropical variety in our paper.

By Theorem 3.4 we are now in the situation to provide an algorithm to find a binomial in an Artinian k-algebra. Let $T = \mathbb{Q}[x_1^{\pm}, \ldots, x_n^{\pm}]$ and $I \subset T$ an Artinian ideal. This means that T/I is a finite-dimensional $\mathbb{Q}$-vector space on which the x_i act by multiplication.[4]

The central facts about this are:

- Multiplication with x_i is a linear map $X_i \in \mathrm{End}(T/I)$.
- Monomials corresponds to compositions.
- The endomorphisms X_i commute pairwise.
- The endomorphisms X_i are all invertible.

Our aim is to test if a binomial $x^u - \lambda x^v$ is in I. But since we are in a Laurent polynomial ring, it suffices to test if $x^u - \lambda \in I$. The same question is, if some composition X^u of the X_i is $\lambda \,\mathrm{id}_{T/I}$, a scalar multiple of the identity. The indetermiantes here are u and λ. It would look much nicer if we could get rid of λ, would it not? It turns out that we can! For this let $l = \dim_{\mathbb{Q}} T/I$ and let K be a finite extension of $\mathbb{Q}$ which contains the l-th roots of the determinants of the X_i. Let $M_i = X_i / \sqrt[l]{\det X_i}$.

Proposition 3.5. *Let $u \in \mathbb{Z}^n$. There exists $\lambda \in K$ such that $x^u - \lambda \in I$ if and only if $M^u = \mathrm{id}_{T/I}$.*

Proof. We are asking for $X^u = \lambda \,\mathrm{id}_{T/I}$, but taking determinants on both sides shows that in this case $\prod_i (\det X_i)^{u_i} = \lambda^l$. So the result follows from the definition of the M_i. □

Now, computing all u such that $M^u = \mathrm{id}_{T/I}$ for commuting matrices M_i over finite extension fields of $\mathbb{Q}$ is a solved problem in number theory. It was solved by Ge, a student of Hendrik W. Lenstra for M_i being just elements of a number field (i.e. 1×1-matrices)[Ge93] and then extended by [BBC+96] to commuting invertible matrices. The lattice of solutions u can be computed in polynomial time. See [LS18] for a modern survey.

Remarks.

- Membership in a subgroup of a matrix group is in general undecidable
- This algorithm is *not* implemented. Anders implemented an efficient computation of $T(I)^{\perp}$, but we did not get any farther.
- There is the much harder problem of deciding if I contains a binomial after a linear change of coordinates on $k[x_1, \ldots, x_n]$. See [KMM19] for deciding if an ideal *is binomial* after a coordinate change.

[4] But we don't want to find short polynomials in here, but in the still infinite vector space I.

4 No short polynomials vanish on all fixed rank matrices

In this lecture we consider ideals generated by the $(r+1)$-minors of matrices. This is based on [DKW21]. Throughout the arguments work over an algebraically closed field, but this makes the situation easier and provides a stronger theorem.

We want to show the following theorem and variants:

Theorem 4.1. *Any polynomial vanishing on all matrices of rank r has at least* $(r+1)!$ *terms.*

Corollary 4.2. *Over any field, there are no polynomials with fewer than* $(r+1)!$ *terms in the ideal generated by all* $r+1$*-minors of a generic matrix.*

Proof. Extending coefficients to algebraic closure might pop up shorter polynomials, but then we have the bound from the theorem. □

Remarks.

- Below we also discuss a version of the theorem for skew-symmetric matrices of fixed rank (on which shorter polynomials vanish).
- Symmetric matrices are an open problem!
- The paper also contains a characterisation of the shortest polynomials. If we have some polynomial f vanishing on X, we can produce more polynomials with the same number of terms and also vanishing on X.. First, any monomial multiple $x^u f$ does. And second, if $\mathrm{char}(k) = p$ is finite, then there is the beginners binomial theorem and f^p has the same number of terms and also vanishes. The effective version of Theorem 4.1 says that monomial times p^e-th power are the only $(r+1)!$-short polynomials vanishing on all rank r-matrices. The proofs of these characterisations are somewhat long and also use Gröbner techniques.

Generic linear spaces

Our approach is based on induction via Laplace expansion of determinants. As a necessary preparation we state a proposition about short polynomials on linear spaces. In fact, this one major technical ingredient in [DKW21].

The polynomial ring $k[x_1, \ldots, x_n]$ consists of all polynomial functions on k^n. With respect to the standard basis x_i extracts the i-th entry of $v \in k^n$. Assume $U \subset k^n$ is a linear space of dimension r. There is always a linear polynomial with $r+1$-terms that vanishes on U. There do exist linear spaces like $U = k^r \times \{0\}$ where even a variable vanishes, but this does not happen „in general“.

The set of all subspaces of dimension r is $\mathrm{Gr}_r(k^n)$, the *Grassmannian*. It is a projective variety (and a manifold, and studied all over mathematics). Not being a special subspace is an algebraic condition on the Grassmannian. This means that *most* subspaces are not special.

A subspace U is *very general* if it avoids countably many algebraic conditions. Over an uncountable field, very general subspaces always exist.

Proposition 4.3. *Any polynomial vanishing on a very general subspace U of dimension r has at least $r+1$ terms.*

Remarks.

- There is also a characterisation: The minimizing polynomials are monomial times p^e-th power of an already vanishing linear form.
- Very general means that there are a countable number of algebraic conditions. The infinite part coming from the countable many degrees that could potentially contribute.
- The proof of the bound uses elimination theory.

Rank r matrices

We consider matrices in $k^{m\times n}$ of rank at most $r \leq m, n$. Matrices with rank at most r are cut out by the vanishing of the $(r+1)$-minors, which have $(r+1)!$ terms.

Theorem 4.4. *Any polynomial vanishing on all rank r matrices in $K^{m\times n}$ has at least $(r+1)!$ terms.*

Proof. The proof is by induction on r. For $r = 0$ the statement is evidently true. Now we suppose that $r \geq 1$ and that the statement is true for $r-1$.

Let f be a nonzero polynomial that vanishes on all rank-r matrices. Then $m, n > r$. Furthermore, since the matrices of rank at most r form an affine cone, any homogeneous component of f also vanishes on them, hence we may assume that f is homogeneous of positive degree.

Let $x_m = (x_{m1}, \ldots, x_{mn})$ be variables representing the last row of the matrix, and write

$$f = \sum_{\alpha\in S} f_\alpha x_m^\alpha.$$

where S is a finite subset of $\mathbb{Z}^n_{\geq 0}$ and the f_α are nonzero polynomials in the entries of the first $m-1$ rows. If some f_α vanishes identically on rank-r matrices, then we replace m by $m-1$ and f by that f_α. After finitely many such steps, we reach a situation where no f_α vanishes identically on rank-r matrices.

Each f_α vanishes on every rank-$(r-1)$ matrix of size $(m-1)\times n$. Indeed, if A is such a matrix, then $f(A, x_m)$ is the zero polynomial because appending any m-th row to A yields a matrix of rank at most r, on which f was assumed to vanish. By the induction assumption, each f_α has at least $(r-1)!$ terms.

On the other hand, since no f_α vanishes on all rank-r matrices, for any sufficiently general $(m-1)\times n$-matrix A of rank r, we have $f_\alpha(A)\neq 0$ for all $\alpha\in S$. Now $f(A, x_m)$ vanishes identically on the r-dimensional row space of A. We may further assume that the row space $U\subseteq K^n$ of A is sufficiently general in the sense of Proposition 4.3. Then, by that proposition, $f(A, x_m)$ has at least $r+1$ terms, and hence f has at least $(r+1)\cdot r! = (r+1)!$ terms. □

Remark. A countable number of exceptional conditions enters, because at the moment we cannot exclude the situation that each degree d forces different sufficiently general conditions.

Rank r skew-symmetric matrices

A variant of the above proof also shows that no short polynomial vanishes on all skew-symmetric matrices. We need the following facts from linear algebra. If $A=-A^T\in K^{n\times n}$ is skew-symmetric, then the rank of A is an even number. In particular, $\det(A)=0$ if n is odd (since $\det(A)=(-1)^n\det(A)$. If n is even, then $\det(A)$ is a square of a polynomial called the *Pfaffian*. An $n\times n$ Pfaffian has much fewer terms than a determinant, namely $n!! := 1\cdot 3\cdots(n-1)$. For example, the (4×4)-Pfaffian is a quadric with $1\cdot 3=3$ and this vansishes on all rank ≤ 3 skew-symmetric matrices of format 4×4.

Theorem 4.5. *Let r be even and let $m\geq r$. There is no nonzero polynomial vanishishing on all skew-symmetric $m\times m$-matrices of rank $\leq r$ that has fewer than $(r+1)!!$ terms.*

Remarks.

- Again, any polynomial with $(r+1)!!$ terms that vanishes on all skew-symmetric $m\times m$-matrices of rank r is a one-term multiple of a p^e-th power of some principal $(r+2)$-Pfaffian, for some $e\in\mathbb{Z}_{\geq 0}$.
- For any field L, the ideal I in the polynomial ring $K[x_{ij}\mid 1\leq i<j\leq m]$ generated by the r-Pfaffians of the matrix x does not contain polynomials with fewer than $(r+1)!!$ terms, and the only polynomials in I with $(r+1)!!$ terms are those in Theorem 4.5.

Proof. The proof proceeds along the same lines as that of Theorem 4.1. Again, we proceed by induction on r. For $r=0$, the 2-Pfaffians are precisely the matrix entries, which of course are the shortest nonzero polynomials vanishing on the zero matrix.

Assume that $r \geq 2$ and decompose

$$f = \sum_{\alpha \in S} f_\alpha x_m^\alpha.$$

where x_m consists of the first $m-1$ entries of the last row—which are, up to a sign, also the first $m-1$ entries of the last column—and where the f_α are nonzero polynomials in the (lower-triangular) entries of the top left $(m-1) \times (m-1)$-block.

Now all f_α vanish on all skew symmetric matrices A of rank at most $r-2$. Indeed, for an arbitrary row vector $u \in K^{m-1}$, the skew symmetric matrix

$$\begin{bmatrix} A & -u^T \\ u & 0 \end{bmatrix}$$

has rank at most r, and hence f vanishes on it. Therefore, if some $f_\alpha(A) \neq 0$, then $f(A, x_m)$ is a nonzero polynomial that vanishes identically on K^{m-1}, a contradiction since K is infinite. From the induction hypothesis, we conclude that each f_α has at least $(r-1)!!$ terms.

We may further assume that no f_α vanishes identically on rank-r skew-symmetric matrices; otherwise, we would replace f by f_α. Pick a sufficiently general skew symmetric matrix $A \in K^{(m-1)\times(m-1)}$ of rank r. Then $f_\alpha(A) \neq 0$ for all α, and we claim that $f(A, x_m)$ is a polynomial that vanishes identically on the row space of A. Indeed, if u is in the row space of A, then appending it to A as an m-th row does not increase the rank of A, and then appending $-u^T$, along with a zero, as the last column, could only increase the rank by 1, but since a skew-symmetric matrix has even rank, it does not. Hence f vanishes on the resulting matrix, and thus $f(A, x_m)$ vanishes on the row space of A. By Proposition 4.3, at least $r+1$ of the f_α are nonzero. Therefore, f has at least $(r+1) \cdot ((r-1)!!) = (r+1)!!$ terms, as desired. □

The proof techniques above cannot be applied in a straightforward way to symmetric matrices because here the rank can grow by 2 when appending a row and column, e.g.

$$\operatorname{rk} \begin{pmatrix} 1 & 1 \\ 1 & 1 \end{pmatrix} = 1 \qquad \text{while} \qquad \operatorname{rk} \begin{pmatrix} 1 & 1 & 1 \\ 1 & 1 & 0 \\ 1 & 0 & 0 \end{pmatrix} = 3.$$

When taking arbitrary minors, corresponding to rows $I \subset [n]$ and columns $J \subset [n]$ of symmetric matrices, different types of polynomials occur, depending on the size of $I \cap J$. For example, if $I = J$ the minor is a determinant of a symmetric matrix. Then the number of terms equals the number of collections of necklaces

that can be made with r distinct beads (see https://oeis.org/A002135), while if $I \cap J = \emptyset$, an r-minor is just a usual determinant with $r!$ terms. When attempting to prove that the shortest polynomials are principal minors using Laplace expansion, all of these different polynomials need to be taken into account in the induction step.

References

[BBC+96] László Babai, Robert Beals, Jin-yi Cai, Gábor Ivanyos, and Eugene M. Luks, *Multiplicative equations over commuting matrices*, Proceedings of the Seventh Annual ACM-SIAM Symposium on Discrete Algorithms (Atlanta, GA, 1996), ACM, New York, 1996, pp. 498–507.

[BG02] Winfried Bruns and Joseph Gubeladze, *Polytopal linear retractions*, Transactions of the American Mathematical Society **354** (2002), no. 1, 179–203.

[DKW21] Jan Draisma, Thomas Kahle, and Finn Wiersig, *No short polynomials vanish on bounded rank matrices*, preprint, arXiv:2112.11764 (2021).

[ES96] David Eisenbud and Bernd Sturmfels, *Binomial ideals*, Duke Mathematical Journal **84** (1996), no. 1, 1–45.

[Ge93] Guoqiang Ge, *Algorithms related to multiplicative representations*, Ph.D. thesis, University of California, Berkeley, 1993.

[GRT10] Mark Giesbrecht, Daniel S Roche, and Hrushikesh Tilak, *Computing sparse multiples of polynomials*, International Symposium on Algorithms and Computation, Springer, 2010, pp. 266–278.

[JKK17] Anders Jensen, Thomas Kahle, and Lukas Katthän, *Finding binomials in polynomial ideals*, Research in the Mathematical Sciences **4** (2017), no. 1, 1–10.

[Kho91] Askold G Khovanskii, *Fewnomials*, vol. 88, American Mathematical Soc., 1991.

[KMM19] Lukas Katthän, Mateusz Michałek, and Ezra Miller, *When is a polynomial ideal binomial after an ambient automorphism?*, Foundations of Computational Mathematics **19** (2019), no. 6, 1363–1385.

[LS18] Hendrik W Lenstra and Alice Silverberg, *Algorithms for commutative algebras over the rational numbers*, Foundations of Computational Mathematics **18** (2018), no. 1, 159–180.

[McC83] S Thomas McCormick, *A combinatorial approach to some sparse matrix problems*, Tech. report, Stanford University, Systems optimization lab, 1983.

[Min76] Edward Minieka, *Finding the circuits of a matroid*, Journal of Research of NIST – B **80B** (1976), no. 3.

[MR19] Diane Maclagan and Felipe Rincón, *Tropical schemes, tropical cycles, and valuated matroids*, Journal of the European Mathematical Society **22** (2019), no. 3, 777–796.