Computing Small Certificates of Inconsistency of Quadratic Fewnomial Systems

Jean-Charles Faugère, Pierre-Jean Spaenlehauer, Jules Svartz

Computing Small Certificates of Inconsistency of Quadratic Fewnomial Systems

Jean-Charles Faugère[1], Pierre-jean Spaenlehauer[2], and Jules Svartz[3]

[1]Sorbonne Universités, UPMC, Univ. Paris 06, CNRS , INRIA, Laboratoire d'Informatique de Paris 6 (LIP6), Équipe PolSys, 4 place Jussieu, 75252 Paris Cedex 05, France. jean-charles.faugere@inria.fr

[2]Inria, CNRS, Université de Lorraine, Nancy, France. pierre-jean.spaenlehauer@inria.fr

[3]Ministère de l'Éducation Nationale, Lycée Masséna, Sorbonne Universités, UPMC, Univ. Paris 06, CNRS , INRIA, Laboratoire d'Informatique de Paris 6 (LIP6), Équipe PolSys, 4 place Jussieu, 75252 Paris Cedex 05, France. jsvartz@ens-cachan.fr

Abstract

Bézout's theorem states that dense generic systems of n multivariate quadratic equations in n variables have 2^n solutions over algebraically closed fields. When only a small subset $\mathbf{M}$ of monomials appear in the equations (*fewnomial systems*), the number of solutions may decrease dramatically. We focus in this work on subsets of quadratic monomials $\mathbf{M}$ such that generic systems with support $\mathbf{M}$ do not admit any solution at all. For these systems, Hilbert's Nullstellensatz ensures the existence of algebraic certificates of inconsistency. However, up to our knowledge all known bounds on the sizes of such certificates —including those which take into account the Newton polytopes of the polynomials— are exponential in n. Our main results show that if the inequality $2|\mathbf{M}| - 2n \leq \sqrt{1+8\nu} - 1$ holds for a quadratic fewnomial system – where ν is the matching number of a graph associated with $\mathbf{M}$, and $|\mathbf{M}|$ is the cardinality of $\mathbf{M}$ – then there exists generically a certificate of inconsistency of linear size (measured as the number of coefficients in the ground field $\mathbb{K}$). Moreover this certificate can be computed within a polynomial number of arithmetic operations. Next, we evaluate how often this inequality holds, and we give evidence that the probability that the inequality is satisfied depends strongly on the number of squares. More precisely, we show that if $\mathbf{M}$ is picked uniformly at random among the subsets of $n+k+1$ quadratic monomials containing at least $\Omega(n^{1/2+\varepsilon})$ squares, then the probability that the inequality holds tends to 1 as n grows. Interestingly, this phenomenon is related with the matching number of random graphs in the Erdös-Renyi model. Finally, we provide experimental results showing that certificates in inconsistency can be computed for systems with more than 10000 variables and equations.

1 Introduction

Context and problem statement. Identifying classes of structured polynomial systems and designing dedicated algorithms to solve them is a central theme in computer algebra and in computational algebraic geometry, due to the wide range of applications where such systems appear. We investigate quadratic systems involving a small number of monomials (*quadratic fewnomial*

systems). Let $\mathbb{K}$ be a field, $\overline{\mathbb{K}}$ its algebraic closure and $\mathbf{M}$ be a finite subset of monomials of degree at most two in a polynomial ring $\mathbb{K}[X_1, \ldots, X_n]$. Suppose also that the constant 1 belongs to $\mathbf{M}$, and let $\mathcal{L}_{\mathbf{M}}$ be the $\mathbb{K}$-linear space of polynomials in $\mathbb{K}[X_1, \ldots, X_n]$ spanned by $\mathbf{M}$.

A classical question is to bound the number of solutions in $\overline{\mathbb{K}}^n$ of a system $f_1(X_1, \ldots, X_n) = \cdots = f_n(X_1, \ldots, X_n) = 0$, where all polynomials lie in $\mathcal{L}_{\mathbf{M}}$ and have generic coefficients. When the exponent vectors of the monomials in $\mathbf{M}$ are the points with integer coordinates in a lattice polytope, Kushnirenko's theorem states that the number of toric solutions (*i.e.* solutions whose all coordinates are nonzero) is bounded by the normalized volume of the polytope [24]. A variant of this theorem indicates that such generic polynomial systems do not admit any solution if the dimension of the $\mathbb{Q}$-linear space generated by the exponent vectors of the monomials in $\mathbf{M}$ does not equal n.

Solving a polynomial system and deciding if it has any solution in $\overline{\mathbb{K}}^n$ are two closely related questions. One classical method to produce a certificate that a polynomial system does not have any solution is to provide an algebraic relation via Hilbert's Nullstellensatz:

Problem 1 - Effective fewnomial Nullstellensatz. Given a system $(f_1, \ldots, f_m) \in \mathcal{L}_{\mathbf{M}}^m$ such that $f_1(X) = \cdots = f_m(X) = 0$ has no solution in $\overline{\mathbb{K}}^n$, compute $h_1, \ldots h_m \in \mathbb{K}[X_1, \ldots, X_n]$ such that $\sum_{i=1}^m f_i\, h_i = 1$.

Bounding the sizes of the polynomials h_i is a crucial question to estimate the complexity of this problem. In this paper, the notion of size that we use is the number of coefficients in $\mathbb{K}$ required to describe the polynomials $h_1, \ldots, h_m$.

On the other hand, the specification of "solving a polynomial system" depends on the context. If the number of solutions in the algebraic closure $\overline{\mathbb{K}}$ is finite, one way to represent them symbolically is to provide a rational parametrization of their coordinates by the roots of a univariate polynomial. For the sake of simplicity, we consider only the problem of computing a univariate polynomial whose roots parametrize the set of solutions:

Problem 2 - Partial 0-dimensional fewnomial system solving. Given a polynomial system $f_1 = \cdots = f_m = 0$ with support $\mathbf{M}$ that have finitely-many solutions in $\overline{\mathbb{K}}^n$ and a monomial $\mu \in \mathbf{M}$, compute a univariate polynomial $P_\mu \in \mathbb{K}[\mu]$ which vanishes at all the solutions of the system.

Hence, the roots of the univariate polynomial P_μ contain the images of the solutions of the input multivariate system via the monomial map $(X_1, \ldots, X_n) \mapsto \mu$.

Related works. Sparse elimination theory for solving systems with special monomial structures have been developed since the 80s [33]. Several lines of work have been initiated during this period. When the exponent vectors of the monomials occurring in the polynomials of the system are the lattice points in a lattice polytope, connections with convex and toric geometry have been established and dedicated solving methods have been designed: homotopy continuation methods [18, 35], resultants [33, 6], Gröbner bases [34, 13], etc. One important theme of these developments is to relate algebraic structures with combinatorial properties of convex bodies. In particular, Kushnirenko and Bernshtein's theorems [24, 2] provide bounds on the number of isolated toric solutions in terms of the Newton polytopes of the input polynomials. Another line of

work have been initiated by Khovanskii in the 80s on fewnomial systems [19]. The main theme in this setting is to relate the algebraic and algorithmic complexity of several problems to the number of monomials occurring in the equations. For instance, a classical and challenging question is to bound the number of real solutions in the positive orthant, see *e.g.* [19, 30, 3, 4, 20, 21] for results on this topic. Bounding the size of a certificate of inconsistency of a polynomial system via the Nullstellensatz is a classical problem. Up to our knowledge, all known upper bounds on the size of such certificates are exponential in the number of variables n; moreover, examples by Masser and Philippon and by Lazard and Mora show that one cannot hope for better bounds in the worst case. A classical bound is given by Kollar [22]: if the maximal degree of the input inconsistent system $f_1, \ldots, f_m$ is D, then there exist $h_1, \ldots, h_m$ such that $\sum_{i=1}^m f_i\, h_i = 1$ and the degrees of the h_i are bounded by $n \min(n,m) D^{\min(n,m)} + \min(n,m) D$. This bound is general and does not require any further assumption. It was later improved to $\deg(f_i h_i) \leq \max(3, D)^n$ [14]. When there is no solution at infinity, the degrees of the polynomials h_i are bounded by $(D-1)n$ [25, 5]: the number of coefficients in dense polynomials of this degree is still exponential in n. For general polynomial systems, it would be surprising that certificates of inconsistency with size polynomial in the input size exist, as this would imply $NP = coNP$. Estimates taking into account the bitsize of the coefficients that appear in the certificate in terms of the bitsize of the coefficients of the input system are provided by *Arithmetic Nullstellensätze*, see *e.g.* [23] and references within. Two milestones on the sparse effective Nullstellensatz are the bounds in [6] and [31]: these bounds provide certificates of size bounded by a quantity which depends on the Newton polytopes of the input polynomials. However, their size is exponential in the size of the input, and these bounds do not take into account the sparsity of the support inside its Newton polytope. One of the main difficulty to generalize these techniques to fewnomial quadratic systems is the fact that the proofs rely on algebraic properties (normality, Cohen-Macaulay algebras) that hold for semigroup algebras generated by lattice points in normal polytopes, but not for semigroup algebras generated by a scattered set of monomials.

Other models of sparse systems have also been investigated. For instance, systems where each quadratic equation involves a small subset of variables have been investigated in [29] and [7]. Connections between combinatorial properties of graphs and polynomial systems is a classical topic which has been investigated from several viewpoints. For instance, square-free monomial ideals have many combinatorial properties and can be seen as the *edge ideals* of graphs, see *e.g.* [15, 17]. Connections between the regularity of the edge ideal of a graph and its matching number and cochordal cover number are shown in [36]. Cohen-Macaulay criteria for such ideals are investigated in [16, 10].

Bounds on the size of certificate of inconsistency of polynomial systems are a important ingredient in algebraic proof complexity, see *e.g.* the Nullstellensatz proof system [1] and related works [8].

Main results. An open question is whether there exist certificates of inconsistency of polynomial size for general fewnomial systems involving $n + k + 1$ monomials in n variables. The goal of this work is to investigate this question in the case of quadratic polynomials. We present an explicit criterion which identifies subsets $\mathbf{M}$ of monomials of degree at most 2 such that systems of n equations in n variables with support $\mathbf{M}$ and generic coefficients do not have any solution, and such that there exists a sparse certificate $\sum_{i=1}^n f_i h_i = 1$, where all polynomials $h_1, \ldots, h_n$ lie in $\mathcal{L}_{\mathbf{M}}$. Therefore, the number of coefficients in $\mathbb{K}$ required to represent the certificate is the same as that of the input system. Moreover, when $\mathbf{M}$ is such a subset, we propose a method which computes such $h_1, \ldots, h_n$ within a polynomial number of arithmetic operations.

More precisely, we model the set $\mathbf{M}$ by a graph G on $n+1$ vertices, where each edge represents a nonconstant square-free monomial in $\mathbf{M}$. The constant 1 and the squares in $\mathbf{M}$ are distinguished with loops in the graph (the precise construction is described in Section 3). Let $\nu(\mathbf{M})$ denote the *matching number* (*i.e.* the maximum cardinality of a matching) of the subgraph of vertices in G with a loop.

Theorem 1.1. *If* $m \geq |\mathbf{M}| - \frac{\sqrt{1+8\nu(\mathbf{M})}-1}{2}$, *then a generic system* $(f_1, \ldots, f_m) \in \mathcal{L}_{\mathbf{M}}^m$ *has no solution in* $\overline{\mathbb{K}}^n$. *Moreover, there exists* $(h_1, \ldots, h_m)$ *solving Problem 1 s.t. all* h_i *lie in* $\mathcal{L}_{\mathbf{M}}$ *and they can be computed within* $O\left(m^\omega \binom{|\mathbf{M}|+1}{2}^\omega\right)$ *operations in* $\mathbb{K}$, *where* ω *is a feasible exponent for matrix multiplication (*$\omega < 2.37286$ *with Le Gall's algorithm [26]).*

We would like to emphasize that the inequality $m \geq |\mathbf{M}| - \frac{\sqrt{1+8\nu(\mathbf{M})}-1}{2}$ can be checked in polynomial time, since the matching number of a graph can be computed in polynomial time with Edmonds' algorithm [11]. Next, we relate how often the assumptions of Theorem 1.1 hold with the number of squares in the support $\mathbf{M}$. If the subset of square monomials and the subset of square-free monomials in $\mathbf{M}$ are chosen at random, and the cardinality of $\mathbf{M}$ is $n+k+1$ and the number of squares is larger than $\Omega(n^{1/2+\varepsilon})$ for some $\varepsilon > 0$, then the assumptions of Theorem 1.1 hold with large probability, leading to the following statement:

Theorem 1.2. *Let* k *be a fixed integer,* $a_n, b_n \in \mathbb{N}$ *be such* $a_n + b_n = n + k + 1$, *and* $\mathbf{M}$ *be a subset of monomials of degree at most* 2 *in* $\mathbb{K}[X_1, \ldots, X_n]$ *distributed uniformly at random among those that contain the constant* 1, a_n *nonsquare monomials, and* $b_n - 1$ *non-constant square monomials. Assume further that* $b_n = \Omega(n^{1/2+\varepsilon})$, *for* $\varepsilon > 0$. *Then the probability that the assumptions of Theorem 1.1 with* $m = n$ *are satisfied for* $\mathbf{M}$ *tends towards* 1 *as* n *grows.*

The cornerstones of the proof of this theorem rely on properties of random graphs in the Erdös-Renyi model. Experiments suggest that this result is sharp: when there are at most $O(n^{1/2})$ square monomials in $\mathbf{M}$, we observe experimentally that the probability of having a certificate in $\mathcal{L}_{\mathbf{M}}^m$ seems to converge to a non-zero value smaller than 1 as n grows. This is also the case when the support is chosen uniformly at random (the expected number of squares is $O(1)$). We propose a conjecture stating that the limit probability is nonzero in that case.

We also study a limit case: when $|\mathbf{M}| = n + k + 1$ and all the squares are in $\mathbf{M}$. The generic number of solutions in this setting is given by the Bézout's theorem: it equals 2^n. We shall see that with probability tending to 1, these solutions can be compactly represented as the orbits of 2^{2k+2} points under an action of $(\mathbb{Z}/2\mathbb{Z})^{n-2k-2}$. Computing these solutions amounts to solving a system of $2k+2$ equations in $2k+2$ variables: the time complexity of this task does not depend on n. A direct consequence is that computing a compact representation of the solutions of such systems require a number of operations in $\mathbb{K}$ which is polynomial in n, even though their number of solutions is exponential in n. This suggests that the number of solutions in the algebraic closure, which is often used to measure the complexity of solving polynomial systems, might in some cases greatly overestimate the complexity for fewnomial systems. Another open issue is to extend this work to the non quadratic case.

Finally, we show experimental results obtained with our proof-of-concept implementation. They show that certificates of inconsistency can be computed for quadratic fewnomial systems with more than 10000 variables and equations when there are sufficiently many squares in the monomial support. Moreover, we also observe some unexpected behaviors which raise new questions about

fewnomial systems. For instance, as n grows, there seems to be a phase transition in the probability of having a small certificate of inconsistency. Moreover, in the case where there are few squares in the fewnomial system (this case is not covered by the theoretical analysis), there seems to be a non-zero probability that a fewnomial system has a small certificate of inconsistency. These phenomenons remain to be explained.

Organization of the paper. Section 2 introduces notation and states preliminary results. The core result of the paper is proved in Section 3, establishing a connection between the matching number and the existence of a small certificate of inconsistency. Section 4 is devoted to a probabilistic analysis of the matching number of some random graphs in the Erdös-Renyi model. Section 5 investigates some families of fewnomial systems where all squares appear in the equations. Finally, we report experimental results in Section 6 and state a conjecture for quadratic fewnomial systems involving few square monomials.

2 Notation and preliminaries

Notation. Throughout this paper, $\mathbb{K}$ denotes a field of odd characteristic. Its algebraic closure is denoted by $\overline{\mathbb{K}}$. If $X_1, \ldots, X_n$ are variables, and $\alpha \in \mathbb{N}^n$, then the shorthand X^α stands for the monomial $X_1^{\alpha_1} \ldots X_n^{\alpha_n}$. The symbol $\mathbf{M}$ denotes a finite subset of monomials in $\mathbb{K}[X_1, \ldots, X_n]$ containing the constant 1. For any $i \in \mathbb{N}$, $\mathbf{M}^i$ denotes the subset of all products of i monomials in $\mathbf{M}$. Its cardinality is denoted by $|\mathbf{M}^i|$. By convention, $|\mathbf{M}^0| = 1$. By slight abuse of notation, we call *dimension* of an ideal I in a ring R the Krull dimension of the quotient ring R/I.

Complexity model. Complexity bounds in this paper count the number of operations $\{+, -, \times, \div\}$ in the field $\mathbb{K}$. It is not our goal to take into account the bitsize of the coefficients in $\mathbb{K}$. Hence, we count each arithmetic operation with unit cost. We do not take into account operations on monomials. The notion of size that we use for polynomial systems is the number of coefficients in $\mathbb{K}$ required to represent them. Note that if $\mathbb{K}$ is a finite field, then the bitsizes of the elements in $\mathbb{K}$ are bounded, and hence the bit complexity is the same as the arithmetic complexity. Given partial functions g, h from a set I to $\mathbb{N}$, we use the following classical Landau notation: $f = O(g)$ means that f/g is bounded above by a constant, $f = \Omega(g)$ is equivalent to $g = O(f)$, and $f = \Theta(g)$ means that $f = O(g)$ and $g = O(f)$.

Genericity. Let $\mathcal{L}_{\mathbf{M}}$ denote the $\mathbb{K}$-linear space spanned by $\mathbf{M}$. It has dimension $|\mathbf{M}|$. We say that a property holds for a generic system $(f_1, \ldots, f_m) \in \mathcal{L}_{\mathbf{M}}^m$ if there exists a dense Zariski open subset $\mathcal{O}$ of $\mathcal{L}_{\mathbf{M}}^m$ s.t. this property holds for any system in $\mathcal{O}$.

Semigroup algebras. The main algebraic structure that we consider are *semigroup algebras* (also called *toric rings*): if $\mathbf{M} \subset \mathbb{K}[X_1, \ldots, X_n]$ is a finite subset of monomials containing 1, we let $\mathbb{K}[\mathbf{M}]$ denote the subalgebra of $\mathbb{K}[X_1, \ldots, X_n]$ generated by $\mathbf{M}$. We do not make any assumption on the Krull dimension of the ring $\mathbb{K}[\mathbf{M}]$. Semigroup algebras which are domains are the coordinate rings of affine toric varieties [9]. We refer to [28, Ch. 7] for a more detailed presentation. By slight abuse of notation, we call variety of a system $f_1, \ldots, f_m \in \mathbb{K}[\mathbf{M}]$ the variety in $\overline{\mathbb{K}}^n$ associated to the ideal $\langle f_1, \ldots, f_n \rangle \subset \mathbb{K}[X_1, \ldots, X_n]$.

The following proposition is a variant of the weak Nullstellensatz for the total coordinate ring of projective toric varieties (see *e.g.* [9, Prop. 5.2.6]).

Proposition 2.1. *The variety associated with a system $f_1, \ldots, f_m \in \mathbb{K}[\mathbf{M}]$ is empty if and only if there exist $h_1, \ldots, h_m \in \mathbb{K}[\mathbf{M}]$ such that $\sum_{i=1}^m f_i h_i = 1$.*

Proof. the ring $\mathbb{K}[\mathbf{M}]$ is isomorphic to $\mathbb{K}[X]/I_M$, where I_M is a toric ideal generated by binomials $b_1, \ldots, b_\ell$. Let $\widetilde{f_1}, \ldots, \widetilde{f_m}$ be the images of $f_1, \ldots, f_m$ by the isomorphism. Using the Nullstellensatz on the system $\widetilde{f_1}, \ldots, \widetilde{f_m}, b_1, \ldots, b_\ell$ in $\mathbb{K}[X]$ and pulling it back to $\mathbb{K}[\mathbf{M}]$ proves the proposition. □

Proposition 2.1 indicates that we can look for polynomial relations in $\mathbb{K}[\mathbf{M}]$ instead of the whole algebra $\mathbb{K}[X]$. Although narrowing the search for the certificate in $\mathbb{K}[\mathbf{M}]$ instead of $\mathbb{K}[X]$ constrains the problem, we shall see that this approach enables us to find efficiently small certificates. This leads to the following variant of Problem 1:

Problem 3 - Effective fewnomial Nullstellensatz in $\mathbb{K}[\mathbf{M}]$. Given a system $f_1, \ldots, f_m \in \mathbb{K}[\mathbf{M}]$ and such that $f_1(X) = \cdots = f_m(X) = 0$ has no solution in $\overline{\mathbb{K}}^n$, compute $h_1, \ldots, h_m \in \mathbb{K}[\mathbf{M}]$ such that $\sum_{i=1}^m f_i\, h_i = 1$.

3 Monomials and support graphs

In this section, we show a connection between graphs and properties of $\mathbb{K}[\mathbf{M}]$. In particular, we focus on quadratic relations between monomials in $\mathbf{M}$, *i.e.* at $\mathbb{K}$-linear relations in the vector space spanned by $\mathbf{M}^2$. We start by adding a new variable X_0 and by considering the homogenized support $\mathbf{M}^h = \{X_0^{2-\deg(\mu)}\mu\}_{\mu\in\mathbf{M}}$. We associate with $\mathbf{M}$ a simple labeled undirected graph G on $S = \{0, \ldots, n\}$ whose edges are $E = \{(i,j) \mid X_iX_j \in \mathbf{M}^h, i \neq j\}$. There is a loop at a vertex i iff $X_i^2 \in \mathbf{M}^h$.

Example 3.1. *Let* $\mathbf{M} = \{1, X_1^2, X_2^2, X_3^2, X_3, X_4, X_1X_2, X_2X_3, X_3X_4\}$. *The following picture represents the graph G; squares in $\mathbf{M}^h$ are indicated by a loop.*

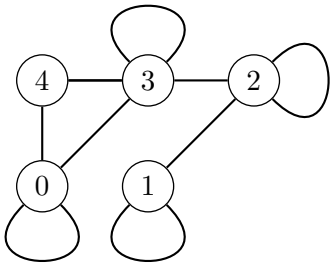

Quadratic relations between elements of $\mathbf{M}$ are of the form $\mu_1\,\mu_2 = \mu_3\,\mu_4$ for (not necessarily distinct) monomials $\mu_1, \mu_2, \mu_3, \mu_4 \in \mathbf{M}$. For the quadratic supports $\mathbf{M}$ that we consider in this paper, these quadratic relations come in three flavors that appear as subgraphs of G and are described in Figure 1. The next proposition shows how the cardinality of $\mathbf{M}^2$ can be computed from the number of quadratic relations and the number of 4-cliques in G. We recall that a 4-clique is a subgraph on 4 vertices such that every pair of vertices is linked by an edge.

Proposition 3.2. *The cardinality of* $\mathbf{M}^2$ *equals* $\binom{|\mathbf{M}|+1}{2} - \lambda(G) + \mathrm{clique}_4(G)$, *where* $\lambda(G)$ *is the number of subgraphs of G isomorphic to any of the three graphs in Figure 1 and* $\mathrm{clique}_4(G)$ *is the number of 4-cliques in G.*

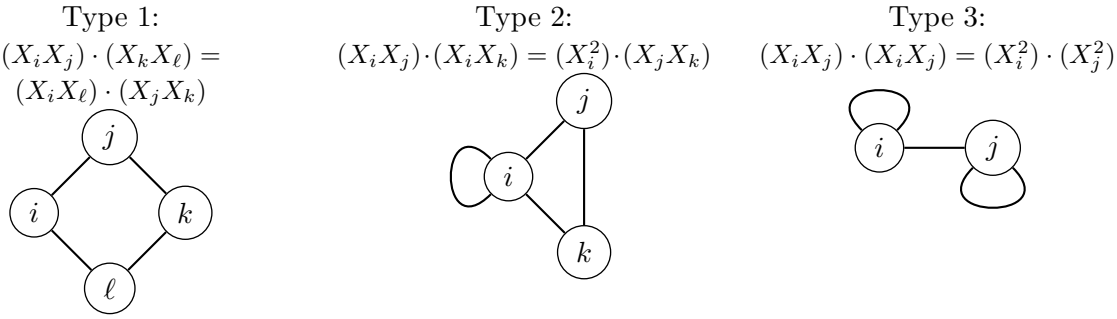

Figure 1: The three types of quadratic relations

Proof. We can form $\binom{|\mathbf{M}|+1}{2}$ products of two (non-necessarily distinct) elements in $\mathbf{M}$. However, some of these products are counted several times because of the quadratic relations between elements in $\mathbf{M}$. This is corrected by the terms $-\lambda(G) + \text{clique}_4(G)$; we detail below the possible cases:

- If $\mu = X_iX_jX_kX_\ell$ is a product of four distinct variables, then μ can be obtained from $\mathbf{M}$ by three different products, since $\mu = (X_iX_j)(X_kX_\ell) = (X_iX_k)(X_jX_\ell) = (X_iX_\ell)(X_jX_k)$. Depending on the number of pairs of such edges that lie in the graph, the monomial μ is counted one, two or three times in $\binom{|\mathbf{M}|+1}{2}$.
  If there is only one way to obtain μ (for example if (X_iX_j) and (X_kX_ℓ) are the only monomials in $\mathbf{M}$ whose products are μ), then the subgraph associated with the vertices $\{X_i, X_j, X_k, X_\ell\}$ is neither of type 1 nor a 4-clique. Hence, μ is counted only one time.
  If there are two ways to obtain μ, then the subgraph associated with the vertices $\{X_i, X_j, X_k, X_\ell\}$ is of type 1 but not a 4-clique. Hence, μ is counted twice in $\binom{|\mathbf{M}|+1}{2}$ but this is corrected by the term $\lambda(G)$.
  If all the three products are possible, then the subgraph associated with the vertices $\{X_i, X_j, X_k, X_\ell\}$ contains three subgraphs of type 1, and is also a 4-clique. Therefore μ is counted 3 times in $\binom{|\mathbf{M}|+1}{2}$, removed 3 times in $\lambda(G)$ and counted once in $\text{clique}_4(G)$.
- If $\mu = X_i^2X_jX_k$ is a monomial involving three distinct variables, then μ is counted twice in $\binom{|\mathbf{M}|+1}{2}$ if and only the subgraph associated with the vertices $\{X_i, X_j, X_k\}$ is of type 2. In this case one contribution is removed by the term $\lambda(G)$, hence μ is counted one time.
- Similarly, monomials $\mu = X_i^2X_j^2$ are counted once or twice in the formula $\binom{|\mathbf{M}|+1}{2}$: if it is counted twice (*i.e.* when $X_iX_j, X_i^2, X_j^2 \in \mathbf{M}$), then the subgraph associated with $\{X_i, X_j\}$ is of type 3. □

Notation. For a graph G associated with a set of monomials $\mathbf{M}$, let G' be the subgraph of squares (*i.e.* the subgraph of vertices with a loop).

Definition 3.3. *A* matching *(also called independent edge set) of G' is a set of edges of G' without common vertices. We let $\nu(\mathbf{M})$ denote the* matching number *of G',* i.e. *the maximum cardinality of a matching of G'.*

The matching number of a graph can be computed in polynomial time by Edmonds's algorithm [11]. We refer to [27] for more details on matching theory. We state now the main result of this section, which connects the matching number of the graph G' to the existence of a small certificate of inconsistency:

Theorem 3.4. *Let $(f_1, \ldots, f_m) \in \mathcal{L}_{\mathbf{M}}^m$ be a system with generic coefficients. If $m \geq |\mathbf{M}| - \frac{\sqrt{1+8\nu(\mathbf{M})}-1}{2}$, then there exist polynomials $h_1, \ldots, h_m \in \mathcal{L}_{\mathbf{M}}$ such that $\sum_{i=1}^m f_i h_i = 1$.*

The proof of this theorem is postponed to the end of the section. It is actually not surprising that systems satisfying the assumptions of Theorem 3.4 do not have any solution, since the dimension of the $\mathbb{Q}$-vector space generated by the exponent vectors in $\mathbf{M}$ is upper bounded by $|\mathbf{M}| - \nu(\mathbf{M})$: each edge (i,j) in G' means that $X_i^2, X_j^2, X_iX_j \in \mathbf{M}$ and the exponent vectors of these three monomials are linearly dependent over $\mathbb{Q}$. The main point of Theorem 3.4 is that, under the condition on $\nu(\mathbf{M})$, the polynomials $(h_1, \ldots, h_m)$ for the effective Nullstellensatz can be searched in $\mathcal{L}_{\mathbf{M}}$. This allows to obtain to get small certificates of inconsistency:

Corollary 3.5. *With the notation and under the assumptions of Theorem 3.4, there is an explicit algorithm which solves Problem 1 within $O\left(m|\mathbf{M}|\left(\binom{|\mathbf{M}|+1}{2} - \lambda(G) + \mathrm{clique}_4(G)\right)^{\omega-1}\right)$ arithmetic operations, where ω is a feasible exponent for matrix multiplication ($\omega < 2.37286$ with Le Gall's algorithm [26]). This complexity is polynomial in the number of coefficients $m|\mathbf{M}|$ of the input system.*

Proof. According to Theorem 3.4, there exist polynomials $h_1, \ldots, h_m$ with support $\mathbf{M}$ s.t. $\sum_i h_i f_i = 1$, by decomposing the polynomials h_i in the monomial basis there exists a relation

$$\sum_i \sum_{\mu \in \mathbf{M}} \alpha_{\mu,i}\, \mu f_i = 1$$

where $\alpha_{\mu,i} \in \mathbb{K}$. Let $V \subset \mathsf{Span}_{\mathbb{K}}(\mathbf{M}^2)$ be the linear space generated by the products $\{\mu\, f_i\}_{\mu \in \mathbf{M}, i \in \{1,\ldots,m\}}$. Consequently, computing the polynomials h_i amounts to solving a linear system over $\mathbb{K}$ with $m\,|\mathbf{M}|$ unknowns and $|\mathbf{M}^2|$ equations. Solving it requires $O(m|\mathbf{M}| \cdot |\mathbf{M}^2|^{\omega-1})$ operations in $\mathbb{K}$ [32, Prop. 2.11]. Proposition 3.2 concludes the proof. □

The sequel of this section is devoted to this proof of Theorem 3.4. The squareroot involved in the formula is a consequence of the following lemma, as the maximal value of p for which $n \geq \binom{n-p+1}{2}$.

Lemma 3.6. *There exist linear forms $\ell_1, \ldots, \ell_p \in \mathbb{K}[X_1, \ldots, X_n]$ such that the ideal*

$$I = \langle X_1^2, \ldots, X_n^2, \ell_1(X_1, \ldots, X_n), \ldots, \ell_p(X_1, \ldots, X_n)\rangle$$

contains all monomials of degree 2 iff $p \geq n - \frac{\sqrt{1+8n}-1}{2}$.

Proof. The vector space of $(n-p)$-variate quadratic forms has dimension $\binom{n-p+1}{2}$. From the inequality $p \geq n - \frac{\sqrt{1+8n}-1}{2}$, we obtain $n \geq \binom{n-p+1}{2}$. This inequality and the fact that any quadratic form can be written as a linear combination of squares of linear forms (since $\mathrm{char}(\mathbb{K}) \neq 2$), implies that there exist $\ell'_1, \ldots, \ell'_n$ such that their squares ${\ell'_1}^2, \ldots, {\ell'_n}^2$ generate the space of $(n-p)$-variate quadratic forms. Then the dimension of the linear space generated by $\ell'_1, \ldots, \ell'_n$ is necessarily maximal and equals $n-p$. Up to permuting the indices, we assume also that $\ell'_1, \ldots, \ell'_{n-p}$ are linearly independent. Hence the ideal $I' = \langle \ell'_1(X_1, \ldots, X_{n-p})^2, \ldots, \ell'_n(X_1, \ldots, X_{n-p})^2, X_{n-p+1}, \ldots, X_n\rangle$ contains all monomials of degree 2. We rewrite I' as

$$I' = \langle \ell''_1(X_1, \ldots, X_n)^2, \ldots, \ell''_n(X_1, \ldots, X_n)^2, X_{n-p+1}, \ldots, X_n\rangle,$$

$$\begin{cases} \ell''_i(X_1,\dots,X_n) = \ell'_i(X_1,\dots,X_{n-p}) \text{ if } 1 \le i \le n-p \\ \ell''_i(X_1,\dots,X_n) = X_i - \ell'_i(X_1,\dots,X_{n-p}) \text{ otherwise.} \end{cases}$$

Note that the linear forms $\ell''_1,\dots,\ell''_n$ are linearly independent by construction. We consider the automorphism θ of $\mathbb{K}[X_1,\dots,X_n]$ defined by $\theta(X_i) = \ell''_i(X_1,\dots,X_n)$, and we set $\ell_i(X_1,\dots,X_n) = \theta^{-1}(X_{n-p+i})$ for $i \in \{1,\dots,p\}$. Therefore I is the inverse image of I' by θ and hence contains all the monomials of degree 2.

It remains to prove the converse statement, *i.e.* that $p < n - \frac{\sqrt{1+8n}-1}{2}$ implies that there do not exist such linear forms $\ell_1,\dots,\ell_p$. This is achieved by a similar argument: if such linear forms existed, then there would exist a set of n generators of the vector space of $(n-p)$-variate quadratic forms. This is not possible if $p < n - \frac{\sqrt{1+8n}-1}{2}$ since this vector space has dimension $\binom{n-p+1}{2}$. □

We can now prove the main theorem of this section:

Proof of Theorem 3.4. We prove a homogeneous version of Theorem 3.4: let $f_1^{(h)},\dots,f_m^{(h)} \in \mathcal{L}^h_{\mathbf{M}} \subset \mathbb{K}[X_0,\dots,X_n]$ be the homogenization of the generic system $f_1,\dots,f_m$. We shall show that any monomial in $(\mathbf{M}^h)^2$ (see the definition of $\mathbf{M}^h$ at the beginning of this section) belongs to the ideal $\langle f_1^{(h)},\dots,f_m^{(h)}\rangle \subset \mathbb{K}[\mathbf{M}^h]$. This will imply that there exist $h_1^{(h)},\dots,h_m^{(h)} \in \mathbb{K}[\mathbf{M}^h]$ such that $\sum_{i=1}^m f_i^{(h)} h_i^{(h)} = X_0^4 \in (\mathbf{M}^h)^2$. Setting $X_0 = 1$ in this equation yields the desired relation.

First, we prove the existence of one system $f_1^{(h)},\dots,f_m^{(h)}$ such that all monomials of $(\mathbf{M}^h)^2$ appear in the ideal $\langle f_1^{(h)},\dots,f_m^{(h)}\rangle$. Throughout this proof, we let $A = \{\{a_1,b_1\},\dots,\{a_{\nu(\mathbf{M})},b_{\nu(\mathbf{M})}\}\} \subset \{0,\dots,n\}^2$ denote a matching of G' of maximum cardinality. We construct a system from A whose polynomials are:

1. all the monomials in $\mathbf{M}^h$ of the form X_iX_j with $i \neq j$;
2. all the monomials in $\mathbf{M}^h$ of the form X_i^2 with i not appearing in A;
3. for each $i \in \{1,\dots,\nu(\mathbf{M})\}$, the polynomial $X_{a_i}^2 - X_{b_i}^2$;
4. the polynomials $\ell_1(X_{a_1}^2,\dots,X_{a_{\nu(\mathbf{M})}}^2),\dots,\ell_p(X_{a_1}^2,\dots,X_{a_{\nu(\mathbf{M})}}^2)$, where the $\nu(\mathbf{M})$-variate linear forms $\ell_1,\dots,\ell_p$ are obtained by replacing n by $\nu(\mathbf{M})$ in Lemma 3.6.

This is a system of $|\mathbf{M}| - \left\lfloor \frac{\sqrt{1+8\nu(\mathbf{M})}-1}{2} \right\rfloor$ polynomials, generating an ideal $I \subset \mathbb{K}[\mathbf{M}^h]$. We claim that all monomials in $(\mathbf{M}^h)^2$ are in the ideal of $\mathbb{K}[\mathbf{M}^h]$ generated by these polynomials:

- every monomial in $(\mathbf{M}^h)^2$ involving at least 3 different variables belongs necessarily to the ideal generated by the monomials X_iX_j with $i \neq j$; the same holds for monomials of the form $X_i^3X_j$ with $i \neq j$;
- next, we look at monomials of the form X_i^4. If i does not appear in A, then by construction X_i^2 is in the ideal. If i is in A, then there exists j such that $i = a_j$ or $i = b_j$. Noticing that $X_{a_j}^4 = X_{a_j}^2(X_{a_j}^2 - X_{b_j}^2) - (X_{a_j}X_{b_j})^2$ or $X_{b_j}^4 = X_{b_j}^2(X_{a_j}^2 - X_{b_j}^2) - (X_{a_j}X_{b_j})^2$ shows that $X_i^4 \in I$.
- finally, we focus on monomials of the form $X_i^2X_j^2$. If i or j do not appear in A, then either X_i^2 or X_j^2 belongs to I. If both i and j appear in A then Lemma 3.6 tells us that $X_i^2X_j^2$ belongs to the ideal generated by $\langle X_{a_1}^4,\dots,X_{a_{\nu(\mathbf{M})}}^4, \ell_1(X_{a_1}^2,\dots,X_{a_{\nu(\mathbf{M})}}^2),\dots,\ell_p(X_{a_1}^2,\dots,X_{a_{\nu(\mathbf{M})}}^2), X_{a_1}^2 - X_{b_1}^2,\dots,X_{a_{\nu(\mathbf{M})}}^2 - X_{b_{\nu(\mathbf{M})}}^2\rangle$.

So far, we have proven that there exists at least one system such that Theorem 3.4 is correct. It remains to prove that this is true for a generic system. To this end, we note that all monomials in $(\mathbf{M}^h)^2$ belongs to $\langle f_1^{(h)}, \dots, f_m^{(h)} \rangle \subset \mathbb{K}[\mathbf{M}^h]$ if and only if $\mathsf{Span}_{\mathbb{K}}(\{\mu\, f_i^{(h)}\}_{\mu \in \mathbf{M}^h, i \in \{1,\dots,m\}}) = \mathsf{Span}_{\mathbb{K}}((\mathbf{M}^h)^2)$. This is an open condition given by the non-vanishing of a product of minors of the matrix recording the coefficients of $\{\mu\, f_i^{(h)}\}_{\mu \in \mathbf{M}, i \in \{1,\dots,m\}}$. Consequently, there exists a Zariski open subset $\mathcal{O} \subset \mathbb{K}[\mathbf{M}^h]^m$ such that Theorem 3.4 holds. This open subset $\mathcal{O}$ is non-empty by the construction above. The proof is concluded by noticing that any non-empty open subset is dense in the Zariski topology. □

4 Random support graphs

In this section, we assume that the support $\mathbf{M}$ is randomly generated and we estimate the probability that the assumptions of Theorem 3.4 are satisfied. Roughly speaking, the aim of this section is to show that the conditions of Theorem 3.4 hold with large probability if n is large enough and if there are sufficiently many squares in $\mathbf{M}$. Let us consider the following variant of the Erdös-Rényi random graph model: for $n \in \mathbb{N}$, we set two probabilities $p_n, q_n \in [0,1]$, and we consider a sequence of random supports $(\mathbf{M}_n)_{n \in \mathbb{N}}$ where

- $\mathbf{M}_n$ is a subset of quadratic monomials of $\mathbb{K}[X_0, \dots, X_n]$;
- each square X_i^2 appears in $\mathbf{M}_n$ independently with probability q_n;
- each monomial of the form $X_i X_j$ with $i \neq j$ appears independently with probability p_n.

The goal is to estimate in which cases the random variable $\nu(\mathbf{M}_n)$ grows sufficiently quickly so that the assumptions of Theorem 3.4 are satisfied asymptotically with large probability. In order to estimate $\nu(\mathbf{M}_n)$, we first forget the meaning of the graph in terms of monomials and count the number of isolated edges in a random graph G in this variant of the Erdös-Renyi model.

Proposition 4.1. *Let G be a random simple graph on $n+1$ vertices. Each vertex has a loop with probability $q \in [0,1]$ and an edge between any two vertices appear with probability $p \in [0,1]$. Let G' be the subgraph obtained by restricting G to the vertices with a loop and $\mathcal{E}$ be the random variable counting the number of isolated edges in G'. Then $\mathcal{E}$ has expected value and variance*

$$\begin{aligned}
\mathbb{E}(\mathcal{E}) =& \binom{n+1}{2} q^2\, p\, (1 - q(1-(1-p)^2))^{n-1}, \\
\mathrm{Var}(\mathcal{E}) =& \mathbb{E}(\mathcal{E}) - \mathbb{E}(\mathcal{E})^2 + \\
& 6 \binom{n+1}{4} q^4\, p^2\, (1-p)^4 (1 - q(1-(1-p)^4))^{n-3}. \qquad (4.1)
\end{aligned}$$

Proof. For each possible edge e between two vertices $i \neq j$, we denote by X_e the random variable taking the value 1 if e is an isolated edge of G', and 0 otherwise. The probability that i and j appear as vertices with loops in G is q^2. Hence, the probability that the edge e lies in G' is $q^2 p$. Moreover, for a given vertex $k \neq i, j$, the probability that k appears in G' and at least one of the edges (i,k) and (j,k) belong to G is $q(1-(1-p)^2)$. There are $n-1$ other vertices than i and j in G, hence X_e follows a Bernoulli law of parameter $q^2 p (1 - q(1-(1-p)^2))^{n-1}$. It follows that

$$\mathbb{E}(\mathcal{E}) = \sum_e \mathbb{E}(X_e) = \binom{n+1}{2} q^2 p (1 - q(1-(1-p)^2))^{n-1}$$

The computation of the variance can be done similarly. □

We now apply the previous proposition in the case where p and q depends on n, and analyze the convergence of $\mathbb{E}(\mathcal{E})$ and $\mathrm{Var}(\mathcal{E})$ as n grows to infinity.

Corollary 4.2. *Let $p_n = \Theta(n^{-1})$ and $q_n = \Theta(n^{\beta})$. With the notation of Proposition 4.1, if $-1/2 < \beta < 0$ then $\mathbb{E}(\mathcal{E}) = \Theta\left(n^{2\beta+1}\right)$ and $\mathrm{Var}(\mathcal{E}) = \Theta\left(n^{2\beta+1}\right)$.*

Proof. First, note that $\log\left[(1-q_n(1-(1-p_n)^2))^{n-1}\right] = -2np_nq_n + O\left(n^{-1}\right)$ since $\beta < 0$. This shows that $\mathbb{E}(\mathcal{E}) = \frac{q_n^2 p_n n^2}{2} e^{-2np_nq_n}(1+O(n^{-1}))$. The claim on the asymptotic behavior of $\mathbb{E}(\mathcal{E})$ follows from $e^{-2np_nq_n} = \Theta(1)$. Next, let λ denote the last summand in Eq. (4.1), namely $\lambda = \mathrm{Var}(\mathcal{E}) - \mathbb{E}(\mathcal{E}) + \mathbb{E}(\mathcal{E})^2$. The asymptotic behavior of λ can be obtained by a similar analysis:

$$\log\left[(1-q_n(1-(1-p_n)^4))^{n-3}\right] = -4np_nq_n + O\left(n^{-1}\right),$$

hence $\lambda = \frac{q_n^4 p_n^2 n^4}{4} e^{-4np_nq_n}(1+O(n^{-1}))$. Notice that $\mathbb{E}(\mathcal{E})^2 = \frac{q_n^4 p_n^2 n^4}{4} e^{-4np_nq_n}(1+O(n^{-1}))$. Consequently, $\mathbb{E}(\mathcal{E})^2 - \lambda = O(n^{4\beta+1})$, since $e^{-4np_nq_n} = \Theta(1)$. Finally, putting all the estimates together, we obtain $\mathrm{Var}(\mathcal{E}) = \Theta(n^{2\beta+1}) + O(n^{4\beta+1}) = \Theta(n^{2\beta+1})$ since $\beta < 0$. □

Finally, we relate the distribution of $\mathcal{E}$ with the probability that the assumptions of Theorem 3.4 hold for fewnomial systems with $|\mathbf{M}| = n+k+1$. If one wants that $\mathbb{E}(|\mathbf{M}_n|) = n+k+1$ for some fixed k and that the expected number of squares is $(n+1)^{1/2+\varepsilon}$, then one has to choose $q_n = (n+1)^{-1/2+\varepsilon}$ and $p_n = (n+k+1-(n+1)q_n)/\binom{n+1}{2}$. The asymptotic expected behavior of the matching number in that case is described by the following statement:

Lemma 4.3. *Let $\mathbf{M}_n$ be a sequence of random supports where each square monomial appears with probability q_n, and each square-free monomial appears with probability p_n. If $p_n = \Theta(n^{-1})$ and $q_n = \Omega(n^{-1/2+\varepsilon})$, with $0 < \varepsilon < 1/2$, then for any $\ell \in \mathbb{N}$, $\mathbf{P}\left(\nu(\mathbf{M}_n) \geq \ell\right)$ tends towards 1 as n grows.*

Proof. Chebyshev's inequality implies that

$$\begin{aligned}\mathbf{P}(\mathcal{E} \leq \mathbb{E}(\mathcal{E})/2) \leq& \mathbf{P}(\mid \mathcal{E} - \mathbb{E}(\mathcal{E}) \mid \geq \mathbb{E}(\mathcal{E})/2) \leq \frac{4\,\mathrm{Var}(\mathcal{E})}{\mathbb{E}(\mathcal{E})^2} \\ =& O\left(n^{-2(-1/2+\varepsilon)-1}\right) = O\left(n^{-2\varepsilon}\right).\end{aligned}$$

Next, notice that $\mathbb{E}(\mathcal{E})/2 = \Theta\left(n^{2\varepsilon}\right)$ by Corollary 4.2. Also, note that $\mathcal{E} \leq \nu(\mathbf{M}_n)$, so that for n sufficiently large, we have $\mathbf{P}\left(\nu(\mathbf{M}_n) \leq n^{\varepsilon}\right) \leq \mathbf{P}\left(\mathcal{E} \leq n^{2\varepsilon}\right) = O\left(n^{-2\varepsilon}\right)$, which tends towards 0 as n grows. □

Next, we show that these estimates also hold for a different model of random monomial supports. For $n \in \mathbb{N}$ and two integers $a, b \in \mathbb{N}$, we consider the random sets $\mathbf{U}_{n,a,b}$ of quadratic monomials in $\mathbb{K}[X_0, \ldots, X_n]$ distributed uniformly at random among those that contain a non-squares and b squares.

Theorem 4.4. *Let k be a fixed integer, $a_n, b_n \in \mathbb{N}$ be such $a_n + b_n = n+k+1$, and $\mathbf{U}_{n,a_n,b_n}$ be a subset of quadratic monomials in $\mathbb{K}[X_0, \ldots, X_n]$ distributed uniformly at random among those that contain a_n non-square monomials and b_n squares. Assume further that $b_n = \Omega(n^{1/2+\varepsilon})$, for $\varepsilon > 0$. Then the probability that the assumptions of Theorem 3.4 with $m = n$ are satisfied for $\mathbf{U}_{n,a_n,b_n}$ tends towards 1 as n grows.*

Proof. The proof of this theorem is technical and is similar to the classical techniques to prove properties of random graphs in the Erdös-Renyi models [12]. Details are provided in the appendix. □

Theorem 1.2 is a direct consequence of Theorem 4.4 and is obtained by dehomogenization.

5 Systems with all the squares

Next, we investigate the special case of fewnomial systems where all the squares X_i^2 belong to $\mathbf{M}$. This corresponds to a limit case of Theorem 4.4: $\varepsilon = 1/2$. In this setting, the Newton polytopes of the polynomials are the same as those of dense quadratic polynomials, hence these systems have generically 2^n solutions in $\overline{\mathbb{K}}^n$. In the sequel of this section, $\mathbf{M}$ is a set of monomials of degree at most 2 in $\mathbb{K}[X_1, \ldots, X_n]$, of cardinality $n + k + 1$, and which contains the constant 1 and all the squares X_i^2. We also assume that $n > 2k$.

We let ℓ denote the number of variables X_i which appear in a square-free monomial in $\mathbf{M}$. Hence $\ell \leq 2k$. For a 0-dimensional system $(f_1, \ldots, f_n) \in \mathcal{L}^n_{\mathbf{M}}$, we let S denote the $n \times (n-\ell)$ matrix which contains the coefficients of the squares X_i^2 such that X_i does not appear in a square-free monomial in $\mathbf{M}$.

Proposition 5.1. *Let $(f_1, \ldots, f_n) \in \mathcal{L}^n_{\mathbf{M}}$ be a 0-dimensional system with support $\mathbf{M}$. Then the system $f_1 = \cdots = f_n = 0$ has at most 2^n solutions in $\overline{\mathbb{K}}^n$. If the matrix S has full rank, then the solutions are the orbits of at most 2^ℓ points under the action of $(\mathbb{Z}/2\mathbb{Z})^{n-\ell}$ given by*

$$\begin{array}{rccc} \chi: & (\mathbb{Z}/2\mathbb{Z})^{n-\ell} \times \overline{\mathbb{K}}^n & \to & \overline{\mathbb{K}}^n \\ & (\mathbf{e}_i, (a_1, \ldots, a_n)) & \mapsto & (a_1, \ldots, -a_{i_j}, \ldots, a_n) \end{array},$$

where the set $\{i_j\}$ is the set of indices such that X_{i_j} does not appear in a square-free monomial in $\mathbf{M}$.

Proof. Up to a permutation of the indices, we can assume w.l.o.g. that $X_1, \ldots, X_{n-\ell}$ are the variables that does not appear in a square-free monomial in $\mathbf{M}$. Since the matrix S is full-rank, we perform Gaussian elimination to remove the squares $X_{i_j}^2$ which do not belong to an edge of the graph. This provides us with an equivalent system of the form

$$\begin{cases} X_1^2 - g_1(X_{n-\ell+1}, \ldots, X_n) = 0 \\ \qquad \vdots \\ X_{n-\ell}^2 - g_{n-\ell}(X_{n-\ell+1}, \ldots, X_n) = 0 \\ h_1(X_{n-\ell+1}, \ldots, X_n) = \cdots = h_\ell(X_{n-\ell+1}, \ldots, X_n) = 0. \end{cases}$$

We end up with a system $(h_1, \ldots, h_\ell)$ of dense homogeneous polynomials in ℓ variables. Note that ℓ is bounded by $2k$, which does not depend on n. Consequently, this system can be solved within a constant number of operations as n grows. By Bézout theorem, this system has at most 2^ℓ solutions. Finally, if $(a_{n-\ell+1}, \ldots, a_n)$ is a solution of $h_1 = \cdots = h_\ell = 0$, then $(\pm\sqrt{g_1(a_{n-\ell+1}, \ldots, a_n)}, \ldots, \pm\sqrt{g_{n-\ell}(a_{n-\ell+1}, \ldots, a_n)}, a_{n-\ell+1}, \ldots, a_n)$ is a solution of the input system. Moreover, all solutions are of this form. □

Therefore, even though the number of solutions of such systems depends exponentially on n, they can be conveniently represented. Moreover, we show next that computing this representation can be achieved within a number of operations in $\mathbb{K}$ which is polynomial in n:

Corollary 5.2. *Let $(f_1, \ldots, f_n) \in \mathcal{L}_{\mathbf{M}}^n$ be polynomials with support $\mathbf{M}$ satisfying the above assumptions ($|\mathbf{M}| = n+k+1$, all squares are in $\mathbf{M}$, S has full-rank) and μ be a square-free monomial. For fixed k, Problem 2 with input $(f_1, \ldots, f_n)$ and μ can be solved within $O(n^\omega)$ arithmetic operations as n grows, where ω is a feasible exponent for matrix multiplication.*

Proof. With the same notations as in the proof of Proposition 5.1, and by noticing that $\langle h_1, \ldots, h_\ell \rangle \cap \mathbb{K}[\mu] = \langle f_1, \ldots, f_n \rangle \cap \mathbb{K}[\mu]$, solving Problem 2 with input $(h_1, \ldots, h_\ell)$ and μ yields a solution to Problem 2 with input $(f_1, \ldots, f_n)$ and μ. Solving Problem 2 with input $h_1, \ldots, h_\ell$ can be achieved within a time complexity which does not depend on n. Consequently, the only complexity that depends on n is the cost of computing the polynomials $h_1, \ldots, h_\ell$. This is done by linear algebra, within $O(n^\omega)$ operations in $\mathbb{K}$. □

6 Experimental results

In this section, we describe experimental results, validating the theoretical results and illustrating their practical relevance. In particular, our prototype implementation of the algorithm in the proof of Corollary 3.5 is able to compute Nullstellensatz' certificates of inconsistency for systems of 30000 equations and 30000 unknowns generated from the uniform model in Theorem 4.4. This may be compared to the practical timings for solving the same problem with dense generic quadratic systems, where 20 unknowns is already a difficult challenge due to the exponential size of the certificates.

Experimental setting. $\mathbb{K}$ is the finite field GF(65521). The experimental procedure depends on parameters n, k and β:

- generate a random support $\mathbf{M}$ of $n+k+1$ monomials of degree at most 2, containing 1 and $\lfloor n^\beta \rfloor$ squares. The subsets of square monomials and non-square monomials are respectively chosen uniformly at random;
- generate a random system of n equations with support $\mathbf{M}$, where all the coefficients are chosen uniformly at random in $\mathbb{K}$;
- return "success" if our implementation returns a relation $1 = \sum_{i=1}^n h_i f_i$, with $h_i \in \mathcal{L}_{\mathbf{M}}$, else return "failure".

By Theorem 4.4, for any choice of parameters $k \in \mathbb{N}$ and $0.5 < \beta < 1$, the probability that "success" is returned should tend towards 1 as n grows.

First, we study the dependence of the asymptotic behavior on the choice of β. To this end, we fix $k = 1$ and we look at the experimental probability of success as n grows. Experimental results are reported in Figure 2. The results are in accordance with Theorem 4.4: when $\beta > 0.5$, the probability that such systems have no solution and that there exists a Nullstellensatz certificate in $\mathcal{L}_{\mathbf{M}}$ seems to tend to 1 as n grows. We also observe that the convergence seems to depend strongly on β: when β becomes close to the limit value 0.5, the speed of convergence seems to decrease. Next, we focus on the dependency on k. We fix $\beta = 0.9$ and let n grow for different values of k. Experiments are reported in Figure 3. Finally, we look at quadratic supports $\mathbf{M}$ of cardinality $n+k+1$ generated uniformly at random without any constraint on the number of squares. This case is not covered by the analysis of this paper and experiments show a different behavior: the probability of success of the algorithm does not seem to tend to 1 as n grows, contrary to the case $\beta > 0.5$. However, this probability seems to converges to a nonzero value.

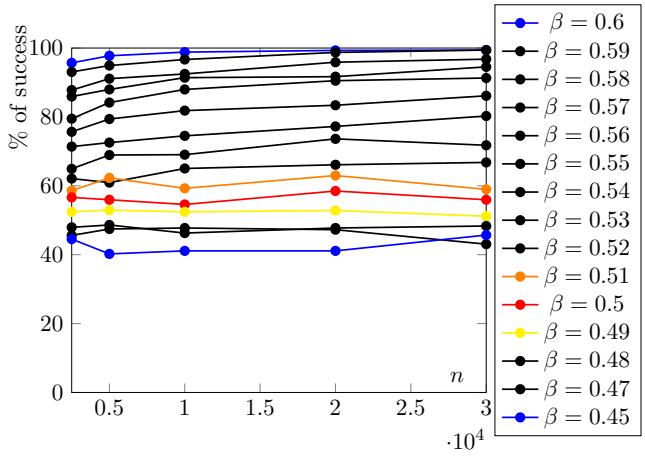

Figure 2: $k = 2$ fixed, n grows, several value of β. Every point is an average over 1000 tests. The relative positions of the curves follow the values of β.

Conjecture 6.1. *Let $k \in \mathbb{N}$ be a fixed integer. For $n \in \mathbb{N}$, let $\mathbf{M}_n$ be a random subset of monomials in n variables of degree at most 2, uniformly distributed among those of cardinality $n + k + 1$ that contain 1. Let $f_1, \ldots, f_n \in \mathcal{L}_{\mathbf{M}}$ be a system with support $\mathbf{M}$ and generic coefficients. Then the probability that there exist $h_1, \ldots, h_n \in \mathcal{L}_{\mathbf{M}}$ such that $\sum_{i=1}^{n} f_i h_i = 1$ tends to a nonzero value as n grows.*

Finally, we report in Figure 4 on experiments about the efficiency our prototype implementation for computing Nullstellensatz certificates. The experiments were conducted on a Mac Retina 2.8Ghz Intel Core i7, and linear algebra computations were performed with `Magma V2.20-3`. We see in these experiments that systems with several thousands of variables can be handled in a few seconds. The algorithm works in two steps: first we reduce the quadratic system with linear algebra (the complexity of this step is independent of β and is represented by the dashed curve); then, the matrix in degree 4 (multiplying all the reduced polynomials by all the monomials in $\mathbf{M}$) is constructed and reduced. The time of this second step depends on β and is indicated by the plain curves. Therefore, these graphs seem to indicate that the cost of computing certificates of inconsistency for these systems is approximately twice the time of computing the row echelon form of a dense $n \times n$ matrix.

References

[1] P. Beame, R. Impagliazzo, J. Krajíček, T. Pitassi, and P. Pudlák. Lower bounds on Hilbert's Nullstellensatz and propositional proofs. In *Foundations of Computer Science, 1994 Proceedings., 35th Annual Symposium on*, pages 794–806. IEEE, 1994.

[2] D. Bernshtein. The number of roots of a system of equations. *Functional Analysis and its Applications*, 9(3):183–185, 1975.

[3] B. Bertrand, F. Bihan, and F. Sottile. Polynomial systems with few real zeroes. *Mathematische Zeitschrift*, 253(2):361–385, 2006.

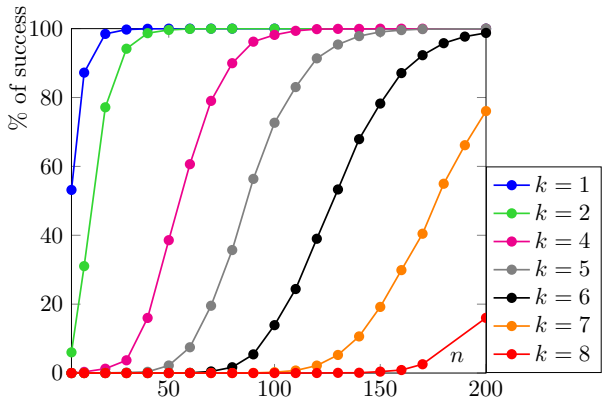

Figure 3: $\beta = 0.9$ fixed, n grows, several value of k. Every point is an average over 10000 tests.

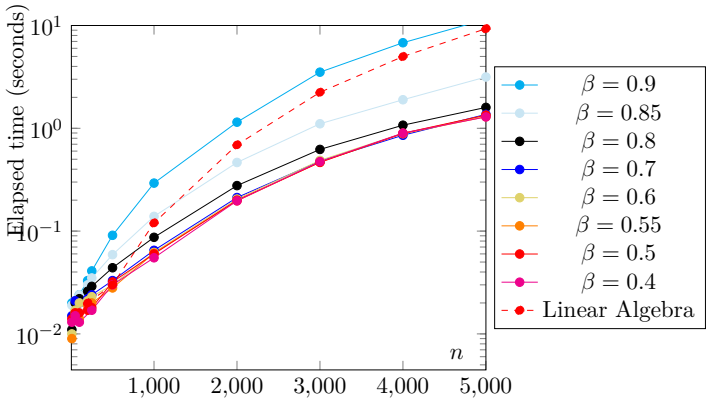

Figure 4: Timings for the computation of the certificate of inconsistency, $k = 2$.

[4] F. Bihan and F. Sottile. New fewnomial upper bounds from Gale dual polynomial systems. *Moscow mathematical journal*, 7(3):387–407, 2007.

[5] W. D. Brownawell. Bounds for the degrees in the Nullstellensatz. *Annals of Mathematics*, pages 577–591, 1987.

[6] J. Canny and I. Emiris. A subdivision-based algorithm for the sparse resultant. *Journal of the ACM*, 47(3):417–451, 2000.

[7] D. Cifuentes and P. Parrilo. Exploiting chordal structure in polynomial ideals: a Gröbner bases approach. *arXiv*, abs/1411.1745, 2014.

[8] M. Clegg, J. Edmonds, and R. Impagliazzo. Using the Groebner basis algorithm to find proofs of unsatisfiability. In *Proceedings of the twenty-eighth annual ACM symposium on Theory of computing*, pages 174–183. ACM, 1996.

[9] D. A. Cox, J. B. Little, and H. K. Schenck. *Toric varieties*. AMS, 2011.

[10] M. Crupi, G. Rinaldo, and N. Terai. Cohen-macaulay edge ideal whose height is half of the number of vertices. *Nagoya Mathematical Journal*, 201:117–131, 2011.

[11] J. Edmonds. Paths, trees, and flowers. *Canadian Journal of mathematics*, 17(3):449–467, 1965.

[12] P. Erdös and A. Rényi. On the evolution of random graphs. *Publ. Math. Inst. Hung. Acad. Sci*, 5:17–61, 1960.

[13] J.-C. Faugère, P.-J. Spaenlehauer, and J. Svartz. Sparse Gröbner bases: the unmixed case. In *Proceedings of ISSAC 2014*, 2014.

[14] N. Fitchas and A. Galligo. Nullstellensatz effectif et conjecture de Serre (théorème de Quillen-Suslin) pour le calcul formel. *Mathematische Nachrichten*, 149(1):231–253, 1990.

[15] R. Fröberg. On Stanley-Reisner rings. *Banach Center Publications*, 26(2):57–70, 1990.

[16] J. Herzog and T. Hibi. Distributive lattices, bipartite graphs and alexander duality. *Journal of Algebraic Combinatorics*, 22(3):289–302, 2005.

[17] J. Herzog, T. Hibi, and X. Zheng. Monomial ideals whose powers have a linear resolution. *Mathematica Scandinavica*, 95(1):23–32, 2004.

[18] B. Huber and B. Sturmfels. A polyhedral method for solving sparse polynomial systems. *Mathematics of Computation*, 64(212):1541–1555, 1995.

[19] A. Khovanskii. On a class of systems of transcendental equations. *Soviet Mathematics Doklady*, 22(3):762–765, 1980.

[20] P. Koiran, N. Portier, and S. Tavenas. On the intersection of a sparse curve and a low-degree curve: A polynomial version of the lost theorem. *Discrete and Computational Geometry*, 53(1):48–63, 2015.

[21] P. Koiran, N. Portier, S. Tavenas, and S. Thomassé. A τ-conjecture for newton polygons. *Foundations of Computational Mathematics*, pages 1–13, 2014.

[22] J. Kollár. Sharp effective Nullstellensatz. *Journal of the American Mathematical Society*, 1(4):963–975, 1988.

[23] T. Krick, L. M. Pardo, M. Sombra, et al. Sharp estimates for the arithmetic Nullstellensatz. *Duke Mathematical Journal*, 109(3):521–598, 2001.

[24] A. G. Kushnirenko. Newton polytopes and the Bezout theorem. *Functional Analysis and its Applications*, 10(3):233–235, 1976.

[25] D. Lazard. Algèbre linéaire sur $k[x_1, \ldots, x_n]$ et élimination. *Bull. Soc. Math. France*, 105:165–190, 1977.

[26] F. Le Gall. Powers of tensors and fast matrix multiplication. In *Proceedings of ISSAC'14*, 2014.

[27] L. Lovász and M. D. Plummer. *Matching theory*. AMS, 1986.

[28] E. Miller and B. Sturmfels. *Combinatorial commutative algebra*, volume 227. Springer Verlag, 2005.

[29] I. Semaev. On solving sparse algebraic equations over finite fields. *Designs, Codes and Crypto.*, 49(1-3):47–60, 2008.

[30] M. Shub and S. Smale. On the intractability of Hilbert's Nullstellensatz and an algebraic version of "P= NP". *Duke Mathematical Journal*, 81(1):47–54, 1996.

[31] M. Sombra. A sparse effective Nullstellensatz. *Advances in Applied Mathematics*, 22(2):271–295, 1999.

[32] A. Storjohann. Algorithms for matrix canonical forms. *Ph.D. thesis*, 2000.

[33] B. Sturmfels. Sparse elimination theory. In *Proceedings of Computational Algebraic Geometry and Commutative Algebra*, pages 377–396. Cambridge Univ. Press, 1991.

[34] B. Sturmfels. *Gröbner bases and convex polytopes*. AMS, 1996.

[35] J. Verschelde, P. Verlinden, and R. Cools. Homotopies exploiting Newton polytopes for solving sparse polynomial systems. *SIAM Journal on Numerical Analysis*, 31(3):915–930, 1994.

[36] R. Woodroofe. Matchings, coverings, and Castelnuovo-Mumford regularity. *J. of Comm. Algebra*, 2:287–304, 2014.

A Proof of Theorem 4.4

Set

$$p_n = a_n / \binom{n+1}{2} \text{ and } q_n = b_n/(n+1),$$

and let $\mathbf{M}_n$ be the random support constructed as above with respect to the probabilities p_n and q_n. We let $\mathbf{MS}_n$ denote the subset of squares in $\mathbf{M}_n$ and $\mathbf{MNS}_n$ denote the subset of nonsquare monomials in $\mathbf{M}_n$. Also, we set $\ell = (k^2+3k+2)/2$. First, we notice that $\mathbf{P}(\nu(\mathbf{M}_n) \geq \ell)$ equals

$$\begin{aligned} &\sum_{\substack{0\leq i\leq \binom{n+1}{2} \\ 0\leq j\leq n}} \mathbf{P}(\nu(\mathbf{M}_n) \geq \ell \mid |\mathbf{MS}| = i \text{ and } |\mathbf{MNS}| = j)\, c_{n,i,j} \\ &\qquad = \sum_{\substack{0\leq i\leq \binom{n+1}{2} \\ 0\leq j\leq n}} \mathbf{P}(\nu(\mathbf{U}_{n,i,j}) \geq \ell)\, c_{n,i,j}, \end{aligned} \tag{A.1}$$

where $c_{n,i,j} = p_n^i (1-p_n)^{\binom{n+1}{2}-i} q_n^j (1-q_n)^{n+1-j} \binom{n+1}{j} \binom{\binom{n+1}{2}}{i}$ is the probability that $|\mathbf{MS}_n| = i$ and $|\mathbf{MNS}_n| = j$. Since the matching number is monotone with respect to the subgraph ordering, we obtain

$$i_1 \geq i_2 \text{ and } j_1 \geq j_2 \Longrightarrow \mathbf{P}(\nu(\mathbf{U}_{n,i_1,j_1}) \geq \ell) \geq \mathbf{P}(\nu(\mathbf{U}_{n,i_2,j_2}) \geq \ell).$$

Consequently, Equation (A.1) implies $\mathbf{P}(\nu(\mathbf{M}_n) \geq \ell)$ is bounded from above by

$$\begin{aligned} &\mathbf{P}(\nu(\mathbf{U}_{n,a_n,b_n}) \geq \ell) \sum_{\substack{0\leq i\leq a_n \\ 0\leq j\leq b_n}} c_{n,i,j} + \\ &\sum_{\substack{0\leq i\leq \binom{n+1}{2} \\ b_n+1\leq j\leq n+1}} \mathbf{P}(\nu(\mathbf{U}_{n,i,j}) \geq \ell)\, c_{n,i,j} + \\ &\sum_{\substack{a_n+1\leq i\leq \binom{n+1}{2} \\ 0\leq j\leq b_n}} \mathbf{P}(\nu(\mathbf{U}_{n,i,j}) \geq \ell)\, c_{n,i,j}. \end{aligned}$$

Note that the first summand is bounded by $\sum_{\substack{0\leq i\leq a_n \\ 0\leq j\leq b_n}} c_{n,i,j}$, the second summand is bounded by $\sum_{\substack{0\leq i\leq \binom{n+1}{2} \\ b_n+1\leq j\leq n+1}} c_{n,i,j}$ and the third one is bounded by $\sum_{\substack{a_n+1\leq i\leq \binom{n+1}{2} \\ 0\leq j\leq b_n}} c_{n,i,j}$. Since the sum of these bounds equals 1, and since the left-hand side of the inequality tends to 1 as n grows by Lemma 4.3, if $\liminf_{n\to\infty} \sum_{\substack{0\leq i\leq a_n \\ 0\leq j\leq b_n}} c_{n,i,j} > 0$, then $\mathbf{P}(\nu(\mathbf{U}_{n,a_n,b_n}) \geq \ell)$ must tend to 1 as n grows.

We prove now that $\liminf_{n\to\infty} \sum_{\substack{0\leq i\leq a_n \\ 0\leq j\leq b_n}} c_{n,i,j} \geq 1/4$. First, we rewrite $\sum_{\substack{0\leq i\leq a_n \\ 0\leq j\leq b_n}} c_{n,i,j}$ as

$$\left(\sum_{0\leq i\leq a_n} p_n^i (1-p_n)^{\binom{n+1}{2}-i} \binom{\binom{n+1}{2}}{i} \right) \left(\sum_{0\leq j\leq b_n} q_n^j (1-q_n)^{n+1-j} \binom{n+1}{j} \right). \tag{A.2}$$

Notice that if $(\mathcal{T}_n)$ is a sequence of random variables following a binomial distribution $B(n, s_n)$ (*i.e.* the sum of n Bernoulli independent variables of parameter s_n) such that $s_n \underset{n\to\infty}{\to} 0$ and $ns_n \underset{n\to\infty}{\to} \infty$, then $(\mathcal{T}_n - ns_n)/\sqrt{ns_n}$ converges in distribution to the standard Gaussian distribution $\mathcal{N}$ (this can be seen on the pointwise convergence of the moment generating function). This implies

$$\liminf_{n\to\infty} \mathbf{P}(\mathcal{T}_n \leq ns_n) \geq \mathbf{P}(\mathcal{N} \leq 0) = 1/2,$$

where $\mathcal{N}$ is a standard Gaussian distribution. Then, we remark that by construction the first factor in Eq. (A.2) equals $\mathbf{P}(|\mathbf{MNS}_n| \leq \binom{n+1}{2} p_n)$ and $|\mathbf{MNS}_n|$ follows a binomial distribution of parameters $(\binom{n+1}{2}, \Omega(1/n))$. Therefore, we obtain

$$\liminf_{n\to\infty} \left(\sum_{0\leq i\leq a_n} p_n^i (1-p_n)^{\binom{n+1}{2}-i} \binom{\binom{n+1}{2}}{i} \right) \geq 1/2.$$

A similar argument shows the same lower bound for the second factor in Eq. (A.2), finishing to prove that $\liminf_{n\to\infty} \sum_{\substack{0\leq i\leq a_n \\ 0\leq j\leq b_n}} c_{n,i,j} \geq 1/4 > 0$.

As explained above, this implies that $\lim_{n\to\infty} \mathbf{P}(\nu(\mathbf{U}_{n,a_n,b_n}) \geq \ell) = 1$ for any $\ell \in \mathbb{N}$. Finally,

$$\begin{aligned} \mathbf{P}(\nu(\mathbf{U}_{n,a_n,b_n}) \geq \ell) &= \mathbf{P}(\nu(\mathbf{U}_{n,a_n,b_n}) \geq (k^2+3k+2)/2) \\ &= \mathbf{P}\left(\frac{\sqrt{1+8\nu(\mathbf{U}_{n,a_n,b_n})}-1}{2} \geq k+1 \right) \\ &\underset{n\to\infty}{\longrightarrow} 1. \end{aligned}$$

This proof is concluded by noticing that $k+1 = |\mathbf{U}_{n,a_n,b_n}| - n$.