Strong Inapproximability for a Promise Rank Problem

Venkatesan Guruswami
Simons Institute for the Theory of Computing, University of California, Berkeley, CA, USA
Departments of EECS and Mathematics, University of California, Berkeley, CA, USA

Xuandi Ren
Department of EECS, University of California, Berkeley, CA, USA

Shaoxuan Tang
Institute for Interdisciplinary Information Sciences, Tsinghua University, Beijing, China

Abstract

Given a linear subspace of $n \times n$ matrices over $\mathbb{F}_{2^r}$ that is promised to contain a matrix of rank 1, we prove that it is hard to find a matrix of rank $n^{o(1/\log\log n)}$, assuming NP doesn't have sub-exponential algorithms. In addition to being a basic problem, the hardness of this problem, even for the exact version, drove recent PCP-free inapproximability results for minimum distance and shortest vector problems concerning codes and lattices.

The proof combines the concept of superposition soundness introduced by Khot and Saket with moment matrices. To produce a rank-gap of 1 vs. k, the reduction runs in time $n^{O(\log k)}$. We also give another moment-matrix-based construction which runs in time $n^{O(k)}$ but works for any finite field $\mathbb{F}_q$.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness; Theory of computation → Approximation algorithms analysis

Keywords and phrases rank minimization, inapproximability, promise problems, PCP, moment matrices

Digital Object Identifier 10.4230/LIPIcs.APPROX/RANDOM.2026.19

Category APPROX

Related Version *Full Version*: `https://arxiv.org/abs/2605.11545`

Funding *Venkatesan Guruswami*: Research supported in part by NSF grant CCF-2211972, a DARPA grant under Contract No. HR0011262E031, and a Simons Investigator award.
Xuandi Ren: Supported in part by NSF grant CCF-2228287 and V.G.'s Simons Investigator award.

Acknowledgements Some of this work was done while Shaoxuan Tang was visiting UC Berkeley. The assistance of ChatGPT was used with some of the proofs in Section 4. All proofs were carefully verified by the authors.

1 Introduction

Given a linear subspace $\mathcal{L} \subseteq \mathbb{F}^{N \times N}$ that is promised to contain a matrix of rank one, we study the problem of finding a nonzero matrix in $\mathcal{L}$ of minimum rank.

This problem is closely related to the problem of finding the *minimum rank distance* of a rank-metric code [5, 9, 17]. In fact, there is a simple and classical embedding of Hamming-metric codes into this setting. Given a linear code $C \subseteq \mathbb{F}^N$, consider the diagonal matrix code

$$\mathrm{Diag}(C) := \{\mathrm{diag}(c) : c \in C\} \subseteq \mathbb{F}^{N \times N}.$$

For every $c \in C$,

$$\operatorname{rank}(\operatorname{diag}(c)) = \|c\|_0,$$

and hence the minimum rank distance of $\operatorname{Diag}(C)$ is exactly the minimum Hamming distance of C. Consequently, hardness results for the Minimum Distance Problem (MDP) [7, 4, 1, 15, 3] directly translate to hardness for approximating the minimum rank in a matrix subspace.

However, this diagonal reduction also explains what it does *not* prove. The YES case inherited from MDP only promises a nonzero matrix of rank $d(C)$, the minimum distance of the code, which is generally not 1. In fact, it is trivial to check if the minimum distance of a code is 1. It is therefore natural to ask whether the problem remains hard under the stronger promise that the subspace contains a nonzero matrix of smallest possible rank, namely rank 1.

A second motivation comes from recent PCP-free inapproximability results for sparse vector problems, including the Minimum Distance Problem and the Shortest Vector Problem [3]. These reductions start from systems of quadratic equations and view each quadratic constraint as a linear constraint on a matrix: if $X = xx^\top$, then a quadratic form in x becomes a linear form in the entries of X. Hence honest solutions to the original quadratic system correspond to rank-one matrices. The main soundness issue is to rule out spurious higher-rank matrices.

In [3], this issue is handled via a non-overlap lemma. Specifically, if C is a linear code of distance $d(C)$ over $\mathbb{F}$, then honest rank-one matrices in the tensor code $C \otimes C$ attain the minimum Hamming weight $d(C)^2$. On the other hand, every matrix in $C \otimes C$ of rank at least 2 has Hamming weight at least $\alpha \cdot d(C)^2$, for some $\alpha > 1$ depending on $\mathbb{F}$. Thus the Hamming-weight gap driving the MDP hardness is, at its core, a structural gap between rank-one matrices and matrices of rank at least two. This suggests stripping away the Hamming-weight objective and asking for the rank gap directly.

In this work, we show strong hardness of distinguishing

$$\exists\, 0 \neq A \in \mathcal{L} \text{ with } \operatorname{rank}(A) = 1 \qquad \text{from} \qquad \forall\, 0 \neq A \in \mathcal{L},\ \operatorname{rank}(A) > k$$

over any fixed finite field of characteristic 2. Specifically, we give a deterministic reduction from 3SAT producing a rank gap of 1 versus k, with matrix dimension $n^{O(\log k)}$. The resulting inapproximability consequences are summarized in the following main theorem.

▶ **Theorem 1.1.** *Let*

$$\operatorname{minrank}(\mathcal{L}) := \min\{\operatorname{rank}(A) : 0 \neq A \in \mathcal{L}\}.$$

For every fixed integer $r \geq 1$, no polynomial-time algorithm can, given a linear subspace $\mathcal{L} \subseteq \mathbb{F}_{2^r}^{N \times N}$, distinguish between

- *(YES)* $\operatorname{minrank}(\mathcal{L}) = 1$,
- *(NO)* $\operatorname{minrank}(\mathcal{L}) > \gamma$,

in the following regimes:

(a) *assuming* $\mathsf{NP} \neq \mathsf{P}$, *when $\gamma > 1$ is any constant;*
(b) *assuming* $\mathsf{NP} \not\subseteq \mathsf{DTIME}(2^{\log^{O(1)} n})$, *when $\gamma = 2^{(\log N)^{1-\epsilon}}$ for any fixed $0 < \epsilon < 1$;*
(c) *assuming* $\mathsf{NP} \not\subseteq \bigcap_{\delta>0} \mathsf{DTIME}(2^{n^\delta})$, *when $\gamma = N^{c/\log\log N}$ for some fixed constant $c > 0$.*

We present the reduction in Section 3. The proof combines the superposition-soundness framework of Khot and Saket [12] with linearized moment matrices. A satisfying assignment to the starting quadratic system gives a rank-one moment matrix. Conversely, a low-rank feasible matrix can be decomposed into a bounded number of symmetric rank-one pieces; these pieces behave like several assignments satisfying the quadratic system in superposition. The soundness analysis of Khot–Saket soundness rules out such superpositions, while the moment-matrix equal-union constraints ensure that the remaining zero-sum case cannot hide a nonzero low-rank matrix.

We also include a simpler direct moment-matrix reduction in Section 4. This construction avoids the superposition-soundness and works over every finite field $\mathbb{F}_q$, at the cost of producing matrices of dimension $n^{O(k)}$. The corresponding inapproximability statement is the following.

▶ **Theorem 1.2.** *Let*

$$\mathrm{minrank}(\mathcal{L}) := \min\{\mathrm{rank}(A) : 0 \neq A \in \mathcal{L}\}.$$

For every fixed finite field $\mathbb{F}_q$, no polynomial-time algorithm can, given a linear subspace $\mathcal{L} \subseteq \mathbb{F}_q^{N \times N}$, distinguish between

- *(YES)* $\mathrm{minrank}(\mathcal{L}) = 1$*,*
- *(NO)* $\mathrm{minrank}(\mathcal{L}) > \gamma$*,*

in the following regimes:

(a) *assuming* $\mathsf{NP} \neq \mathsf{P}$*, when* $\gamma > 1$ *is any constant;*
(b) *assuming* $\mathsf{NP} \not\subseteq \mathsf{DTIME}(2^{\log^{O(1)} n})$*, when* $\gamma = (\log N)^{1-\epsilon}$ *for any fixed* $0 < \epsilon < 1$*;*
(c) *assuming* $\mathsf{NP} \not\subseteq \bigcap_{\delta>0} \mathsf{DTIME}(2^{n^\delta})$*, when* $\gamma = c \log N/ \log\log N$ *for some fixed constant* $c > 0$.

Proof overview

The key ingredient in both reductions is the pseudo-moment matrix. Fix a moment level d. We introduce a pseudo-moment coordinate y_R for every Boolean monomial $x^R := \prod_{i\in R} x_i$ of degree at most $2d$. Note that the vector (y_R) is not assumed to come from an actual Boolean assignment, for which reason it is called pseudo-moment vector. The associated degree-d pseudo-moment matrix[1] is defined by

$$H_d(y)_{S,T} = y_{S\cup T}. \qquad (|S|, |T| \leq d)$$

In the matrix formulation, we impose the equal-union constraints

$$A_{S,T} = A_{S',T'} \qquad \text{whenever} \qquad S \cup T = S' \cup T'.$$

This way, the matrix is not an arbitrary linearization of the products $x^S x^T$. All factorizations of the same Boolean monomial $x^{S\cup T}$ are forced to share a single pseudo-moment coordinate.

This hidden redundancy is what makes rank a useful test. If a pseudo-moment y_R is nonzero, it is not isolated in one entry of the matrix: it appears in every entry whose row and column labels union to R. Therefore a feasible low-rank matrix is constrained not only by the original equations, but also by the truncated Boolean monomial algebra encoded by these equal-union identities.

[1] This is the standard moment-matrix object from the Sum-of-Squares literature: for a pseudoexpectation $\widetilde{\mathbb{E}}$ over the Boolean cube, the moment matrix has entries $\widetilde{\mathbb{E}}[x^S x^T] = \widetilde{\mathbb{E}}[x^{S\cup T}]$. The Lasserre/SoS hierarchy additionally imposes normalization and positive-semidefiniteness of moment matrix. See, e.g., [13, 16, 14, 2].

We first record the simple but powerful observation related to the rank of pseudo-moment matrices that is used in both reductions, and then give an overview of the two reductions separately.

▶ **Lemma 1.3** (Informal structural lemma, used in Lemma 3.10 and Lemma 4.3)**.** *Let A be a matrix satisfying the above equal-union constraints, and let y be the pseudo-moment vector it represents. Suppose R is a minimum-size set with $y_R \neq 0$. Then,*

$$\operatorname{rank}(A) \geq \binom{|R|}{\lfloor |R|/2 \rfloor}.$$

The proof idea is to look at the submatrix with rows indexed by the $\lfloor |R|/2 \rfloor$-subsets $F \subseteq R$ and columns indexed by the $\lceil |R|/2 \rceil$-subsets $G \subseteq R$. Its (F, G) entry is $y_{F \cup G}$. By the minimality of R, this entry is zero unless $F \cup G = R$. Hence the submatrix is a permutation matrix, giving the rank lower bound.

The two reductions use this observation in different ways. In the $n^{O(\log k)}$ superposition reduction, the observation turns the zero-sum case from Khot–Saket soundness into the conclusion that the whole low-rank matrix is zero. In the $n^{O(k)}$ direct pseudo-moment reduction, the same observation forces the existence of a nonzero flat level from which one can round to an honest Boolean solution.

The $n^{O(\log k)}$ superposition reduction over $\mathbb{F}_{2^r}$

We first prove the rank gap over $\mathbb{F}_2$ and then extend it to $\mathbb{F}_{2^r}$ in a black-box way. We adapt the Khot–Saket superposition reduction [12] to produce a constant-free quadratic system $\mathcal{Q}$. After introducing a homogenizing variable x_0, and introducing variables y_S for non-constant monomials x^S of degree at most $d = \Theta(\log k)$, we obtain a quadratic system with the following property. In the YES case, an honest satisfying assignment gives a solution to $\mathcal{Q}$. In the NO case, if t assignments satisfy $\mathcal{Q}$ in superposition, meaning that the sum of their evaluations on every equation of $\mathcal{Q}$ is zero, then their coordinate-wise sum is the zero assignment.

After that, for each quadratic equation

$$\sum_{S,T} c_{S,T} y_S y_T + \sum_R b_R y_R = 0$$

in $\mathcal{Q}$, we linearize it as

$$\sum_{S,T} c_{S,T} A_{S,T} + \sum_R b_R A_{R,R} = 0,$$

and we also impose all equal-union constraints.

For completeness, an honest satisfying assignment gives a feasible rank-one moment matrix.

For soundness, suppose in the NO case that a nonzero feasible matrix A has rank at most k. A decomposition lemma[2] shows

$$A = \sum_{i=1}^{t} u^{(i)} u^{(i)\top}, \qquad t \leq \left\lfloor \frac{3k}{2} \right\rfloor.$$

[2] This decomposition is not obvious: we require a sum of *symmetric* rank-1 terms, which is more restrictive than a general rank decomposition.

Substituting this decomposition into the linearized constraints shows that the assignments $u^{(i)}$ satisfy $\mathcal{Q}$ in superposition[3]. By the soundness guarantee, the sum of these assignments is zero. Consequently, every diagonal entry of A vanishes:

$$A_{S,S} = \sum_i (u_S^{(i)})^2 = \sum_i u_S^{(i)} = 0.$$

The equal-union constraints then imply $A_{S,T} = A_{S\cup T, S\cup T} = 0$ whenever $|S \cup T| \le d$.

We still need to rule out nonzero entries whose union has size larger than d. Suppose such an entry exists, and choose $A_{S,T} \neq 0$ with $|S \cup T|$ minimum. Put $R = S \cup T$. By the structural lemma,

$$\operatorname{rank}(A) \ge \binom{|R|}{\lfloor |R|/2 \rfloor} \ge \binom{d+1}{\lfloor (d+1)/2 \rfloor} > k,$$

contradicting the rank assumption.

Finally, to work over $\mathbb{F}_{2^r}$, we interpret the same $\mathbb{F}_2$-linear equations over the extension field; a nonzero rank-at-most-k matrix over $\mathbb{F}_{2^r}$ would descend, via an $\mathbb{F}_2$-linear functional, to a nonzero matrix over $\mathbb{F}_2$ of rank at most rk. So we just run the base-field reduction with gap rank parameter rk.

The $n^{O(k)}$ direct reduction. This construction starts from QuadEq over an arbitrary finite field $\mathbb{F}_q$. We set $d = k$ and build the pseudo-moment matrix whose rows and columns are indexed by all monomials of degree at most d, including the empty monomial. We impose the same equal-union constraints, and for each equation of the QuadEq instance

$$f_\ell(x) = \sum_{U \subseteq [n],\ |U| \le 2} c_{\ell,U} x^U$$

we impose the following pseudo-moment versions of the identities $x^W f_\ell(x) = 0$:

$$\sum_{\substack{U \subseteq [n] \\ |U| \le 2}} c_{\ell,U} y_{U \cup W} = 0 \qquad \forall W,\ |W| \le 2d - 2\ .$$

These are the pseudo-moment versions of the identities $x^W f_\ell(x) = 0$. The completeness is immediate: a Boolean solution gives a rank-one moment matrix solution, and it is nonzero because the empty-coordinate entry is 1.

For soundness, suppose the source instance is unsatisfiable but there is a nonzero feasible pseudo-moment matrix $H_d(y)$ of rank at most d. Let $r_e = \operatorname{rank} H_e(y)$; this sequence is nondecreasing in e. The structural lemma shows that the first nonzero level already has a reasonably large rank that the sequence r_e cannot keep increasing strictly up to level d. Hence there must be a nonzero flat level

$$r_e = r_{e+1} > 0$$

for some $e < d$.

[3] The constant-free assumption is needed here. If a quadratic equation has a nonzero constant term, substituting a decomposition $A = \sum_{i=1}^t u^{(i)} u^{(i)\top}$ counts that constant with a parity depending on the number of summands, so a linearized feasible matrix yields the desired superposition condition only when t is odd. Huang [11], building on the Khot–Saket superposition framework, uses precisely this oddness phenomenon: he proves that superposition satisfaction by an odd number of assignments is equivalent to odd-covering.

At such a flat level, multiplying a column label by x_i creates no new column direction. Thus multiplication by x_i defines a linear operator T_i on the column space C_e of $H_e(y)$. These operators satisfy the Boolean rules

$$T_i^2 = T_i, \qquad T_iT_j = T_jT_i,$$

and the localizing constraints translate into the operator identities $f_\ell(T) = 0$ on C_e for every source equation f_ℓ. Since the T_i are commuting projections over $\mathbb{F}_q$, they have a common eigenvector with eigenvalues $a_i \in \{0, 1\}$. The tuple $a = (a_1, \dots, a_n)$ is therefore a Boolean point, and the identities $f_\ell(T) = 0$ imply $f_\ell(a) = 0$ for all ℓ, contradicting unsatisfiability.

2 Preliminaries

We now formally define the concepts related to monomial-indexing and pseudo-moments used in Section 3 and Section 4, which we already sketched in the introduction.

Let $[n] := \{1, 2, \dots, n\}$ for any positive integer n. For a subset S, we write $x^S := \prod_{i\in S} x_i$, and we set $x^\emptyset := 1$. All products of monomials and polynomials are taken in the Boolean squarefree monomial algebra. Equivalently, over the relevant field we identify polynomials modulo the relations $x_i^2 = x_i$, so that $x^S x^T = x^{S\cup T}$.

Write

$$\mathcal{U}_{n,d} := \{S \subseteq \{0, 1, \dots, n\} : 1 \le |S| \le d\};$$
$$\mathcal{V}_{n,d} := \{S \subseteq \{1, \dots, n\} : 0 \le |S| \le d\}.$$

Here $\mathcal{U}_{n,d}$ is used for the homogenized $n^{O(\log k)}$ construction: it allows the coordinate 0 but excludes the empty monomial. The set $\mathcal{V}_{n,d}$ is used for the direct $n^{O(k)}$ construction: it uses only the original variables and includes the empty monomial.

When $\mathcal{W}_d$ is either $\mathcal{U}_{n,d}$ or $\mathcal{V}_{n,d}$, define a *monomial assignment of degree d* to be a map

$$\sigma : \{x^S : S \in \mathcal{W}_d\} \to \mathbb{F}.$$

Equivalently, any monomial assignment of degree d can be written as a *degree-d pseudo-moment assignment*, or *pseudo-moment vector*,

$$y = (y_R)_{R\in\mathcal{W}_d}.$$

When a degree-$2d$ pseudo-moment vector is given, its associated *pseudo-moment matrix* is defined as

$$H_d(y) := \left(y_{S\cup T}\right)_{S,T\in\mathcal{W}_d}.$$

If a is an actual Boolean point in $\{0, 1\}^{n+1}$ (or $\{0, 1\}^n$, respectively), the "honest" moment assignment generated by a is $y_R = a^R$. The corresponding moment vector is

$$v_d(a) := (a^S)_{S\in\mathcal{W}_d},$$

and in this case

$$H_d(y) = v_d(a)v_d(a)^\top.$$

In particular, $H_d(y)$ has rank one whenever $v_d(a) \neq 0$.

On the other hand, any matrix $A \in \mathbb{F}^{\mathcal{W}_d\times\mathcal{W}_d}$ satisfying

$$A_{S,T} = A_{S',T'} \qquad \text{whenever } S \cup T = S' \cup T',$$

is a pseudo-moment matrix $H_d(y)$ for some unique pseudo-moment vector $y \in \mathbb{F}^{\mathcal{W}_{2d}}$. We therefore sometimes state the linear constraints on the pseudo-moment matrix in terms of the associated pseudo-moment vector y.

3 An $n^{O(\log k)}$ Construction over $\mathbb{F}_{2^r}$

In this section, we prove the following theorem.

▶ **Theorem 3.1.** *For every fixed integer $r \geq 1$ and every integer $k \geq 1$, there is a deterministic $n^{O(\log k)}$-time reduction which maps a 3SAT instance on n variables to a linear subspace*

$$\mathcal{L} \subseteq \mathbb{F}_{2^r}^{N \times N}, \qquad N = n^{O(\log k)},$$

such that:

- *(YES) if the input formula is satisfiable, then $\mathcal{L}$ contains a nonzero matrix of rank 1;*
- *(NO) if the input formula is unsatisfiable, then $\mathcal{L}$ contains no nonzero matrix of rank at most k.*

To prove Theorem 3.1, we first prove the result over $\mathbb{F}_2$, then transfer it to any $\mathbb{F}_{2^r}$ by a rank-descent argument. Our reduction has four steps:

1. First, we reduce 3SAT to a degree-d polynomial equation system $\mathcal{B}$ over $\mathbb{F}_2$ on the variables $x_0, x_1, \ldots, x_n$. The homogenizing variable x_0 replaces the constant 1 and makes every equation in $\mathcal{B}$ have zero constant term.
2. Next, we replace each nonconstant monomial x^S, $S \in \mathcal{U}_{n,d}$, by a new variable y_S, and we add the moment constraints $y_S y_T = y_{S \cup T}$. This produces a quadratic equation system $\mathcal{Q}$ with zero constant term. Furthermore, we obtain Khot–Saket superposition soundness: in the NO case, if t assignments $\tau^{(1)}, \ldots, \tau^{(t)}$ satisfy every equation of $\mathcal{Q}$ in superposition (see Definition 3.5), then the sum of the t assignments

$$\tau := \sum_{i=1}^{t} \tau^{(i)}$$

must vanish on every y-coordinate.
3. We then linearize the quadratic system $\mathcal{Q}$ by introducing matrix variables $A_{S,T}$ for the products $y_S y_T$. By adding equal-union constraints to enforce the Boolean monomial identity $x^S x^T = x^{S \cup T}$, and by the rank observation in Lemma 1.3, we get the rank gap over $\mathbb{F}_2$.
4. Finally, we extend the hardness from $\mathbb{F}_2$ to $\mathbb{F}_{2^r}$ by interpreting the same homogeneous linear constraints over the extension field. The rank-descent lemma Lemma 3.12 converts any nonzero matrix of rank at most k over $\mathbb{F}_{2^r}$ into a nonzero matrix over $\mathbb{F}_2$ of rank at most rk. Running the $\mathbb{F}_2$ construction with rank parameter rk gives the claimed gap over $\mathbb{F}_{2^r}$.

The four steps are carried out in Sections 3.1–3.4, respectively. For orientation, the first three variable-changing steps are summarized in Figure 1; the corresponding variables and their roles are listed in Table 1.

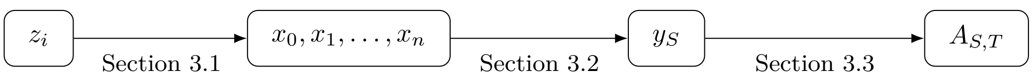

Figure 1 Roadmap of the reduction steps in this section.

Table 1 Variables defined in this section and their roles.

Variables	Meaning	Used in
z_i	original 3SAT variables	starting NP-hardness source
$x_0, x_1, \ldots, x_n$	homogenized Boolean variables	degree-d polynomial system $\mathcal{B}$ with no constant term; x_0 replaces 1
y_S	variable for monomial x^S	quadratic system $\mathcal{Q}$ with Khot–Saket superposition soundness
$A_{S,T}$	matrix variable linearizing $y_S y_T$	linear matrix subspace $\mathcal{L}_d(\mathcal{Q})$ for the rank-gap reduction

3.1 Reducing from 3Sat to constant-free polynomial equations

In this subsection, before presenting our main reduction in Theorem 3.4, we first introduce the following duality fact used in [12]: any assignment to all nonconstant monomials of degree at most d can be represented as the sum of evaluations at actual points of $\mathbb{F}_2^{n+1}$.

▶ **Lemma 3.2** (Monomial assignments as sums of points, [12]). *Let $\mathbb{F}_2[x]_{\leq d}$ denote the $\mathbb{F}_2$-vector space of multilinear polynomials in variables $x_0, \ldots, x_n$ of degree at most d. For every monomial assignment*

$$\sigma : \{x^S : S \in \mathcal{U}_{n,d}\} \to \mathbb{F}_2,$$

define $\sigma(1) = 1$ and extend σ by linearity to all polynomials in $\mathbb{F}_2[x]_{\leq d}$. Then, there is a subset $\beta \subseteq \mathbb{F}_2^{n+1}$ such that

$$\sigma(q) = \sum_{a \in \beta} q(a)$$

holds for every $q \in \mathbb{F}_2[x]_{\leq d}$.

Since $\sigma(1) = 1$, any set β representing σ must have odd cardinality. Among all sets representing σ in this sense, fix one with minimum cardinality and denote it by β_σ. Due to the page limit, we defer the proofs of Lemma 3.2 and the following point-isolator lemma, Lemma 3.3, to the full version of this paper.

▶ **Lemma 3.3** (Point isolator). *Let $\rho \geq 0$ be an integer, $T \subseteq \mathbb{F}_2^{n+1}$ with $|T| < 2^\rho$, and let $a \in T$. Then there exists a multilinear polynomial q of degree at most ρ such that*

$$q(a) = 1 \qquad \textit{and} \qquad q(b) = 0 \quad \textit{for all } b \in T \setminus \{a\}.$$

We now give the reduction from 3SAT to a system of degree-d constant-free polynomial equations. This is a constant-free analogue of the Khot–Saket construction. We homogenize the usual Boolean constraints by introducing a variable x_0, which plays the role of the constant 1 in honest assignments, and by adding the equations $x_i(x_i + x_0) = 0$. If $x_0 = 1$, the clause polynomials encode the original 3SAT instance. If $x_0 = 0$, these equations force every variable to vanish as well.

▶ **Theorem 3.4** (Constant-free low-weight soundness). *For every integer $d \geq 4$, there is a deterministic $n^{O(d)}$-time reduction from a 3SAT instance on n variables to a system $\mathcal{B}$ of constant-free degree-d polynomial equations in $n+1$ variables over $\mathbb{F}_2$ such that:*

- *(YES) if the input formula is satisfiable, then $\mathcal{B}$ has a satisfying assignment $a \in \mathbb{F}_2^{n+1}$ with $a_0 = 1$;*
- *(NO) if the input formula is unsatisfiable and σ is a monomial assignment of degree d satisfying all equations in $\mathcal{B}$, then either $\beta_\sigma = \{\mathbf{0}\}$ or $|\beta_\sigma| \geq 2^{d-3}$.*

Proof. Let the input formula be φ with variables $z_1, \dots, z_n$ and clauses $C_1, \dots, C_m$. We introduce variables $x_0, x_1, \dots, x_n$. For a literal ℓ, define its false-literal form by

$$h_\ell(x) = \begin{cases} x_0 + x_i & \text{if } \ell = z_i, \\ x_i & \text{if } \ell = \neg z_i. \end{cases}$$

When $x_0 = 1$, this is 1 if and only if the literal ℓ is false. For a clause $C_j = \ell_{j,1} \vee \ell_{j,2} \vee \ell_{j,3}$, define the clause polynomial

$$p_j(x) := h_{\ell_{j,1}}(x) h_{\ell_{j,2}}(x) h_{\ell_{j,3}}(x).$$

For every original variable, define the Booleanity polynomial

$$b_i(x) := x_i(x_i + x_0).$$

Every p_j and b_i has no constant term, and their degrees are at most 3 and 2, respectively.

The system $\mathcal{B}$ consists of the following equations:

(1) $x^S p_j(x) = 0$ for every $j \in [m]$ and every monomial x^S, including $S = \emptyset$, of degree at most $d - 3$;

(2) $x^S b_i(x) = 0$ for every $i \in [n]$ and every monomial x^S, including $S = \emptyset$, of degree at most $d - 2$.

All equations have degree at most d and no constant term.

In the YES case, take a satisfying Boolean assignment $(a_1, \dots, a_n)$ to φ and set $a_0 = 1$. Then each $b_i(a)$ vanishes. In every clause at least one literal is true, so at least one of the three corresponding false-literal forms is zero, and hence every $p_j(a)$ vanishes. All monomial multiples in $\mathcal{B}$ therefore vanish at a as well.

Now assume φ is unsatisfiable, and let σ be a monomial assignment satisfying all equations in $\mathcal{B}$. Suppose, for contradiction, that $\beta_\sigma \neq \{\mathbf{0}\}$ and $|\beta_\sigma| < 2^{d-3}$. By the parity observation above, β_σ is nonempty; since it is not $\{\mathbf{0}\}$, it contains a nonzero point. Choose such a point $a \in \beta_\sigma$. We use the following point-isolating lemma to distinguish a from the remaining points in β_σ.

By Lemma 3.3, there exists a polynomial q of degree at most $d - 3$ such that

$$q(a) = 1 \qquad \text{and} \qquad q(b) = 0 \quad \text{for all } b \in \beta_\sigma \setminus \{a\}.$$

First we show that $a_0 = 1$. For every $i \in [n]$, the polynomial $q(x)b_i(x)$ is an $\mathbb{F}_2$-linear combination of equations in $\mathcal{B}$, so

$$0 = \sigma\big(q(x)b_i(x)\big) = \sum_{b \in \beta_\sigma} q(b)b_i(b) = a_i(a_i + a_0). \tag{1}$$

If $a_0 = 0$, then (1) forces $a_i = 0$ for every $i \in [n]$, contradicting the choice of the nonzero point a. Hence $a_0 = 1$.

With $a_0 = 1$, (1) says that the point a obeys the intended Booleanity constraints on the original coordinates, and the values $h_\ell(a)$ are the usual false-literal indicators of the Boolean assignment $(a_1, \dots, a_n)$. Since φ is unsatisfiable, this assignment falsifies some clause, say C_j, and therefore $p_j(a) = 1$. Again $q(x)p_j(x)$ is an $\mathbb{F}_2$-linear combination of equations in $\mathcal{B}$, and hence

$$0 = \sigma\big(q(x)p_j(x)\big) = \sum_{b\in\beta_\sigma} q(b)p_j(b) = q(a)p_j(a) = 1,$$

a contradiction. Thus either $\beta_\sigma = \{\mathbf{0}\}$ or $|\beta_\sigma| \geq 2^{d-3}$. ◀

3.2 Obtaining Khot–Saket superposition soundness

We now define superposition satisfaction, a notion ruled out by the soundness analysis of Khot–Saket. We are particularly interested in this notion because assignments in superposition satisfaction naturally correspond to the factors in the symmetric rank-one decomposition given by Lemma 3.8.

▶ **Definition 3.5** (Superposition satisfaction and aggregate assignment). *Let*

$$g(y) = \sum_{U,V} c_{U,V}\, y_U y_V + \sum_W d_W y_W$$

be a polynomial of degree at most 2 *and without constant term over* $\mathbb{F}_2$*. We say that assignments* $\tau^{(1)}, \dots, \tau^{(t)}$ satisfy $g(y) = 0$ in superposition *if*

$$\sum_{i=1}^{t} g\big(\tau^{(i)}\big) = 0.$$

Their aggregate assignment *is the coordinate-wise sum*

$$\tau := \sum_{i=1}^{t} \tau^{(i)}.$$

For $\beta \subseteq \mathbb{F}_2^{n+1}$ and a polynomial f, write

$$\chi_\beta(f) := (-1)^{\sum_{a\in\beta} f(a)}.$$

The following correlation estimate is from [12, Lemma 3.3], which further builds on [6, Theorem 20].

▶ **Lemma 3.6** (Khot–Saket correlation bound; [12, Lemma 3.3]). *Let* $d \geq 8$ *be a positive multiple of* 4*. Let* σ *be a degree-*d *monomial assignment, and let* β_σ *be a minimum-cardinality set representing* σ *as in Lemma 3.2. Suppose* $|\beta_\sigma| \geq 2^{d-3}$*, and let* $\gamma, \alpha \subseteq \mathbb{F}_2^{n+1}$ *be arbitrary. Let* g *be a uniformly random multilinear polynomial of degree at most* $3d/4$*, and let* h *be a uniformly random multilinear polynomial of degree at most* $d/4$ *with zero constant term. Then*

$$\left|\mathop{\mathbb{E}}_{g,h}\big[\chi_{\beta_\sigma}(gh)\chi_\gamma(g)\chi_\alpha(h)\big]\right| \leq 2^{-2^{d/4-2}+1}.$$

▶ **Theorem 3.7** (Constant-free exact-superposition soundness, adapted from [12]). *There is a constant* $c > 0$ *such that the following holds. Let* $t \geq 1$*, and let* d *be a positive multiple of* 4 *with* $d \geq \max\{8, c\log(t+1)\}$*. From a* 3SAT *instance on* n *variables, one can construct in time* $n^{O(d)}$*, a system* $\mathcal{Q}$ *of quadratic equations with zero constant term, in variables*

$$\{y_S : S \in \mathcal{U}_{n,d}\}$$

with the following properties:

- *(YES) if the* 3Sat *instance is satisfiable, then $\mathcal{Q}$ has a satisfying assignment y with $y_{\{0\}} = 1$;*
- *(NO) if the* 3Sat *instance is unsatisfiable and $\tau^{(1)}, \ldots, \tau^{(t)}$ satisfy all equations of $\mathcal{Q}$ in superposition, then their aggregate assignment $\tau := \sum_{i=1}^{t} \tau^{(i)}$ vanishes on every coordinate y_S, $S \in \mathcal{U}_{n,d}$.*

Moreover, $\mathcal{Q}$ is obtained by replacing each nonconstant monomial x^S in the system $\mathcal{B}$ from Theorem 3.4 with the variable y_S, and by adjoining the constraints

$$y_S y_T = y_{S\cup T} \qquad \textit{for all } S, T \in \mathcal{U}_{n,d} \textit{ with } |S \cup T| \le d.$$

Proof. Apply Theorem 3.4 to obtain the constant-free system $\mathcal{B}$ in the variables $x_0, x_1, \ldots, x_n$. For every $S \in \mathcal{U}_{n,d}$ introduce a variable y_S, intended to represent the monomial x^S. Since every equation of $\mathcal{B}$ has zero constant term, no variable corresponding to $x^{\emptyset}$ is needed. The system $\mathcal{Q}$ contains:

(i) for every equation $\sum_S c_S x^S = 0$ in $\mathcal{B}$, the linear equation $\sum_S c_S y_S = 0$ obtained by replacing each monomial x^S with y_S;

(ii) every multiplicativity constraint $y_S y_T = y_{S\cup T}$ with $S, T \in \mathcal{U}_{n,d}$ and $|S \cup T| \le d$.

The YES case follows directly from Theorem 3.4. If a is the satisfying assignment there, with $a_0 = 1$, then setting $y_S := a^S$ satisfies all linearized equations from $\mathcal{B}$ and all multiplicativity constraints. In particular $y_{\{0\}} = 1$.

For the NO case, let $\tau^{(1)}, \ldots, \tau^{(t)}$ be assignments to the variables y_S satisfying all equations of $\mathcal{Q}$ in superposition. Write

$$\tau_S^{(i)} := \tau^{(i)}(y_S), \qquad \tau_S := \sum_{i=1}^{t} \tau_S^{(i)} \qquad (S \in \mathcal{U}_{n,d}).$$

Thus $\tau = (\tau_S)_{S\in\mathcal{U}_{n,d}}$ is the aggregate y-assignment. Define the associated degree-d monomial assignments σ and $\sigma^{(i)}$ by

$$\sigma(x^S) := \tau_S, \qquad \sigma^{(i)}(x^S) := \tau_S^{(i)} \qquad (S \in \mathcal{U}_{n,d}),$$

extending them to polynomials by the constant-preserving convention from Lemma 3.2. Thus τ lives on the y-coordinates, while σ is the corresponding assignment to the x-monomials.

The main idea here is that, since the equations inherited from $\mathcal{B}$ are linear and homogeneous, superposition satisfaction forces the aggregate assignment $\tau = \sum_i \tau^{(i)}$ to satisfy these equations. Viewing the y-coordinates of τ as a monomial assignment σ, Theorem 3.4 then gives a dichotomy: either $\beta_\sigma = \{\mathbf{0}\}$, or β_σ is large. The latter possibility is ruled out by converting the superposition satisfaction of product equations $y_S y_T = y_{S\cup T}$ into a low-degree correlation identity, and applying the Khot–Saket correlation bound (Lemma 3.6).

The equations inherited from $\mathcal{B}$ become linear after the monomial-to-variable replacement in (i), and they have no constant term. Therefore superposition satisfaction implies that τ satisfies every such linearized equation. Equivalently, σ satisfies every equation of $\mathcal{B}$. By Theorem 3.4, either $\beta_\sigma = \{\mathbf{0}\}$ or $|\beta_\sigma| \ge 2^{d-3}$. If $\beta_\sigma = \{\mathbf{0}\}$, then for every $S \in \mathcal{U}_{n,d}$,

$$\tau_S = \sigma(x^S) = x^S(\mathbf{0}) = 0,$$

which is exactly the claimed vanishing of every coordinate y_S.

It remains to rule out the case $|\beta_\sigma| \ge 2^{d-3}$. For each i, let $\beta_i := \beta_{\sigma^{(i)}}$. The multiplicativity constraint $y_S y_T = y_{S\cup T}$, satisfied in superposition, gives

$$\sigma(x^{S\cup T}) = \tau_{S\cup T} = \sum_{i=1}^{t} \tau_S^{(i)} \tau_T^{(i)} = \sum_{i=1}^{t} \sigma^{(i)}(x^S)\sigma^{(i)}(x^T) \tag{2}$$

whenever $S, T \in \mathcal{U}_{n,d}$ and $|S \cup T| \le d$. Using the same constant-preserving extension, (2) implies, by checking monomial pairs and extending $\mathbb{F}_2$-bilinearly, that

$$\sigma(gh) = \sum_{i=1}^{t} \sigma^{(i)}(g)\sigma^{(i)}(h) \tag{3}$$

for every polynomial g of degree at most $3d/4$ and every polynomial h of degree at most $d/4$ with zero constant term. Indeed, since h has zero constant term, it suffices to check $h = x^T$ and $g = 1$ or $g = x^S$. The case $g = x^S$ is exactly (2); the case $g = 1$ is the coordinate identity $\tau_T = \sum_i \tau_T^{(i)}$. The degree bounds ensure $|S \cup T| \le d$ in the first case.

To analyze (3) using Fourier analysis, we need to switch from $\mathbb{F}_2$ to the real values in $\{-1, 1\}$. For bits $u, v \in \mathbb{F}_2$, if $a = (-1)^u$ and $b = (-1)^v$, then

$$(-1)^{uv} = a \wedge b := \frac{1 + a + b - ab}{2}.$$

Taking signs in (3) and using Lemma 3.2, we obtain, for every such pair g, h,

$$\chi_{\beta_\sigma}(gh) \prod_{i=1}^{t} (\chi_{\beta_i}(g) \wedge \chi_{\beta_i}(h)) = 1. \tag{4}$$

Now choose g and h uniformly at random from the two polynomial spaces above. The expectation of the left-hand side of (4) is 1. On the other hand, expanding each factor

$$a \wedge b = \frac{1 + a + b - ab}{2}$$

expresses the same expectation as a sum of at most 4^t terms, each with coefficient of absolute value at most 2^{-t} and each of the form

$$\mathop{\mathbb{E}}_{g,h}\left[\chi_{\beta_\sigma}(gh)\chi_\gamma(g)\chi_\alpha(h)\right]$$

for some subsets $\gamma, \alpha \subseteq \mathbb{F}_2^{n+1}$. By Lemma 3.6 above, the absolute value of the whole expectation is at most

$$2^t \cdot 2^{-2^{d/4-2}+1}.$$

Choosing the absolute constant c sufficiently large makes this quantity strictly smaller than 1 for every $t \ge 1$, contradicting (4). Hence the case $|\beta_\sigma| \ge 2^{d-3}$ cannot occur, and the aggregate assignment must vanish on every coordinate y_S. ◀

3.3 Linearizing the superposition instance

Fix a degree parameter d, and let $\mathcal{Q}$ be the quadratic system from Theorem 3.7. Its variables are $\{y_S : S \in \mathcal{U}_{n,d}\}$. We define a linear subspace of matrices with rows and columns indexed by $\mathcal{U}_{n,d}$ by imposing two kinds of linear constraints.

(1) *Equal-union constraints:*

$$A_{S,T} = A_{S',T'} \qquad \text{whenever } S \cup T = S' \cup T'. \tag{5}$$

(2) *Equation constraints:* for every equation of $\mathcal{Q}$ in the form

$$\sum_{S,T \in \mathcal{U}_{n,d}} c_{S,T} y_S y_T + \sum_{R \in \mathcal{U}_{n,d}} b_R y_R = 0, \tag{6}$$

where $c_{S,T}, b_R \in \mathbb{F}_2$, impose

$$\sum_{S,T \in \mathcal{U}_{n,d}} c_{S,T} A_{S,T} + \sum_{R \in \mathcal{U}_{n,d}} b_R A_{R,R} = 0. \tag{7}$$

Let $\mathcal{L}_d(\mathcal{Q})$ be the set of matrices satisfying (5) and (7). This is a linear subspace of $\mathbb{F}_2^{N\times N}$, where

$$N = |\mathcal{U}_{n,d}| = \sum_{j=1}^{d} \binom{n+1}{j}.$$

Furthermore, $\mathcal{L}_d(\mathcal{Q})$ only contains symmetric matrices due to the constraints $A_{S,T} = A_{T,S}$ in (5).

▶ **Lemma 3.8** (Decomposition into symmetric rank-one matrices, [12]). *Let $A \in \mathbb{F}_2^{N\times N}$ be symmetric of rank k. Then there is an integer t with $0 \le t \le \lfloor 3k/2 \rfloor$ and vectors $u^{(1)}, \ldots, u^{(t)} \in \mathbb{F}_2^N$ such that*

$$A = \sum_{i=1}^{t} u^{(i)} u^{(i)\top}.$$

A proof of Lemma 3.8 can be found in the full version of this paper.

▶ **Lemma 3.9** (Low rank gives superposition assignments). *Let $k \ge 1$. Let $A \in \mathcal{L}_d(\mathcal{Q})$ have rank at most k, and let $t \ge \lfloor 3k/2 \rfloor$. Then there exist $u^{(1)}, \ldots, u^{(t)} \in \mathbb{F}_2^N$ such that*

$$A = \sum_{i=1}^{t} u^{(i)} u^{(i)\top},$$

and the assignments $y_S \mapsto u_S^{(i)}$ satisfy every equation of $\mathcal{Q}$ in superposition.

Proof. Since $A \in \mathcal{L}_d(\mathcal{Q})$ satisfies the equal-union constraints, it is symmetric. By Lemma 3.8, A is a sum of at most $\lfloor 3k/2 \rfloor \le t$ symmetric rank-one matrices. Padding with zero vectors if necessary, write

$$A = \sum_{i=1}^{t} u^{(i)} u^{(i)\top}.$$

Fix an equation of $\mathcal{Q}$ written as in (6). Since A satisfies the corresponding equation constraint (7), we have

$$0 = \sum_{S,T \in \mathcal{U}_{n,d}} c_{S,T} A_{S,T} + \sum_{R \in \mathcal{U}_{n,d}} b_R A_{R,R}.$$

Substituting the rank-one decomposition and using $(u_R^{(i)})^2 = u_R^{(i)}$ over $\mathbb{F}_2$ gives

$$\begin{aligned} 0 &= \sum_{i=1}^{t} \left(\sum_{S,T \in \mathcal{U}_{n,d}} c_{S,T} u_S^{(i)} u_T^{(i)} + \sum_{R \in \mathcal{U}_{n,d}} b_R (u_R^{(i)})^2 \right) \\ &= \sum_{i=1}^{t} \left(\sum_{S,T \in \mathcal{U}_{n,d}} c_{S,T} u_S^{(i)} u_T^{(i)} + \sum_{R \in \mathcal{U}_{n,d}} b_R u_R^{(i)} \right). \end{aligned}$$

The last expression is the superposition sum of this equation evaluated on the assignments $y_S \mapsto u_S^{(i)}$. Hence the assignments satisfy every equation of $\mathcal{Q}$ in superposition. ◀

▶ **Lemma 3.10** (Zero aggregate assignment forces vanishing). *Let $A \in \mathcal{L}_d(\mathcal{Q})$, $A \neq 0$, admit a decomposition*

$$A = \sum_{i=1}^{t} u^{(i)} u^{(i)\top}$$

such that its aggregate assignment is zero:

$$\sum_{i=1}^{t} u_U^{(i)} = 0 \qquad \textit{for every } U \in \mathcal{U}_{n,d}.$$

Then

$$\operatorname{rank}(A) \geq \binom{d+1}{\lfloor (d+1)/2 \rfloor}.$$

Proof. The zero-aggregation condition gives

$$A_{U,U} = \sum_{i=1}^{t} (u_U^{(i)})^2 = \sum_{i=1}^{t} u_U^{(i)} = 0, \qquad \forall U \in \mathcal{U}_{n,d}.$$

If $U, V \in \mathcal{U}_{n,d}$ and $|U \cup V| \leq d$, then $U \cup V \in \mathcal{U}_{n,d}$, and the equal-union constraints give

$$A_{U,V} = A_{U\cup V, U\cup V} = 0. \tag{8}$$

Suppose $A \neq 0$, and choose U, V with $A_{U,V} = 1$ minimizing $s := |U \cup V|$. By (8), we have $d < s \leq 2d$. Put $R := U \cup V$. For every pair $U', V' \in \mathcal{U}_{n,d}$ with $U' \cup V' = R$, the equal-union constraints give $A_{U',V'} = A_{U,V} = 1$.

Let

$$\mathcal{F} := \{F \subseteq R : |F| = \lfloor s/2 \rfloor\}, \qquad \mathcal{G} := \{G \subseteq R : |G| = \lceil s/2 \rceil\}.$$

Because $d < s \leq 2d$, every set in $\mathcal{F} \cup \mathcal{G}$ lies in $\mathcal{U}_{n,d}$. In the submatrix $A|_{\mathcal{F},\mathcal{G}}$, the entry indexed by (F, G) is 1 exactly when $G = R \setminus F$: in that case $F \cup G = R$, while all other pairs have union a proper subset of R and hence have entry 0 by the minimality of s. Thus $A|_{\mathcal{F},\mathcal{G}}$ is a permutation matrix, so

$$\operatorname{rank}(A) \geq \binom{s}{\lfloor s/2 \rfloor} \geq \binom{d+1}{\lfloor (d+1)/2 \rfloor}. \qquad ◀$$

▶ **Theorem 3.11** (Rank gap over $\mathbb{F}_2$). *For every integer $k \geq 1$, there is a deterministic $n^{O(\log k)}$-time reduction from a* 3SAT *instance on n variables to a linear subspace*

$$\mathcal{L} \subseteq \mathbb{F}_2^{N \times N}, \qquad N = n^{O(\log k)},$$

such that satisfiable instances yield a nonzero rank-1 matrix in $\mathcal{L}$, and unsatisfiable instances yield no nonzero matrix in $\mathcal{L}$ of rank at most k.

Proof. Set

$$t := \left\lfloor \frac{3k}{2} \right\rfloor,$$

and choose d to be a positive multiple of 4 such that

$$d \geq \max\{8, c \log(t+1)\} \qquad \text{and} \qquad \binom{d+1}{\lfloor (d+1)/2 \rfloor} > k,$$

where c is the constant from Theorem 3.7. Then $d = O(\log k)$.

Apply Theorem 3.7 with this d and output $\mathcal{L} := \mathcal{L}_d(\mathcal{Q})$. The construction time is $n^{O(d)} = n^{O(\log k)}$, and

$$N = |\mathcal{U}_{n,d}| = \sum_{j=1}^{d} \binom{n+1}{j} = n^{O(d)} = n^{O(\log k)}.$$

In the YES case, take a satisfying Boolean assignment to the original formula, set $a_0 = 1$, and define $y_S := a^S$ for every $S \in \mathcal{U}_{n,2d}$. Then $H_d(y) = v_d(a)v_d(a)^\top$ is nonzero of rank 1. Since every equation of $\mathcal{Q}$ vanishes at $(a^S)_{S\in\mathcal{U}_{n,d}}$, the matrix $H_d(y)$ satisfies both the equal-union constraints and the equation constraints. Hence $H_d(y) \in \mathcal{L}$.

In the NO case, suppose $A \in \mathcal{L}$ is nonzero and $\operatorname{rank}(A) \leq k$. By Lemma 3.9, there are exactly t assignments whose rank-one sum is A and which satisfy all equations of $\mathcal{Q}$ in superposition. By Theorem 3.7, their aggregate assignment is zero. Then Lemma 3.10 and the choice of d imply that $\operatorname{rank}(A) > k$. ◀

3.4 Extending the hardness from $\mathbb{F}_2$ to $\mathbb{F}_{2^r}$

The following descent lemma is the black-box field-extension step. It loses a factor of r in the rank parameter, which is why the base-field gap below is run with parameter rk. Due to the page limit, we defer its proof to the full version of this paper.

▶ **Lemma 3.12** (Rank descent). *Let $\mathbb{K} = \mathbb{F}_{2^r}$, and let $\mathcal{L} \subseteq \mathbb{F}_2^{N\times N}$ be a linear subspace. If $A \in \mathcal{L} \otimes_{\mathbb{F}_2} \mathbb{K}$ is nonzero and $\operatorname{rank}_{\mathbb{K}}(A) \leq k$, then there exists a nonzero $B \in \mathcal{L}$ with*

$$\operatorname{rank}_{\mathbb{F}_2}(B) \leq rk.$$

Proof of Theorem 3.1. Let $\mathbb{K} = \mathbb{F}_{2^r}$ and set $k_0 := rk$. Apply Theorem 3.11 with rank parameter k_0 to obtain $\mathcal{M} \subseteq \mathbb{F}_2^{N\times N}$, and output

$$\mathcal{L} := \mathcal{M} \otimes_{\mathbb{F}_2} \mathbb{K} \subseteq \mathbb{K}^{N\times N},$$

equivalently the same homogeneous linear equations interpreted over $\mathbb{K}$.

Completeness is preserved under field extension: a nonzero rank-1 matrix in $\mathcal{M}$ remains a nonzero rank-1 matrix in $\mathcal{L}$. For soundness, suppose that the input formula is unsatisfiable and that $A \in \mathcal{L}$ is nonzero with $\operatorname{rank}_{\mathbb{K}}(A) \leq k$. By Lemma 3.12, there is a nonzero $B \in \mathcal{M}$ with $\operatorname{rank}_{\mathbb{F}_2}(B) \leq rk = k_0$, contradicting the soundness of $\mathcal{M}$.

The running time and dimension are

$$n^{O(\log k_0)} = n^{O(\log(rk))}, \qquad N = n^{O(\log k_0)} = n^{O(\log(rk))}.$$

For fixed extension degree r, this is $n^{O(\log k)}$ time and $N = n^{O(\log k)}$. ◀

The inapproximability statement Theorem 1.1 is a canonical corollary of Theorem 3.1.

▶ Remark 3.13 (Toward arbitrary finite fields). It is natural to ask whether the $n^{O(\log k)}$ rank-gap reduction of this section can be obtained over an arbitrary fixed finite field $\mathbb{F}_q$. The main technical step to achieve this is replacing the binary low-degree long code used in the Dinur–Guruswami and Khot–Saket frameworks [6, 12] with a suitable q-ary analogue. In particular, carrying out the superposition-soundness argument over $\mathbb{F}_q$ requires extending the relevant low-error Reed–Muller testing and, crucially, generalizing the correlation bound of Lemma 3.6 to arbitrary finite fields.

4 An $n^{O(k)}$ Construction over any $\mathbb{F}_q$

In this section, we start with the canonical NP-hardness of Boolean QuadEq, and prove Theorem 4.2 by a direct moment-matrix construction.

▶ **Lemma 4.1** (Boolean QuadEq hardness, [8, 10]). *Let $\mathbb{F}_q$ be any finite field. Given squarefree quadratic polynomials $f_1, \ldots, f_m$ over $\mathbb{F}_q$ in variables $x_1, \ldots, x_n$, it is NP-hard to decide whether there exists $a \in \{0,1\}^n \subseteq \mathbb{F}_q^n$ such that $f_1(a) = \cdots = f_m(a) = 0$. Equivalently, each input polynomial may be written as*

$$f_\ell(x) = \sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} x^U, \qquad c_{\ell,U} \in \mathbb{F}_q,$$

where the coordinate indexed by $\emptyset$ is the constant term.

▶ **Theorem 4.2.** *Fix a finite field $\mathbb{F}_q$. For every integer $k \geq 1$, there is a deterministic $n^{O(k)}$-time reduction from a Boolean* QuadEq *instance with n variables over $\mathbb{F}_q$ to a linear subspace*

$$\mathcal{L} \subseteq \mathbb{F}_q^{N \times N}, \qquad N = n^{O(k)},$$

such that:

- *(YES) if the Boolean* QuadEq *instance is satisfiable, then $\mathcal{L}$ contains a nonzero matrix of rank 1;*
- *(NO) if the Boolean* QuadEq *instance is unsatisfiable, then $\mathcal{L}$ contains no nonzero matrix of rank at most k.*

Throughout this section all vector spaces, matrices, and polynomials are over a fixed finite field $\mathbb{F}_q$. We use the Boolean monomial and pseudo-moment conventions from Section 2; in particular, $\mathcal{V}_{n,d}$ includes the empty set.

4.1 The pseudo-moment subspace

Fix an integer $d \geq 1$. Following the notation of Section 2, a degree-$2d$ pseudo-moment vector is $y = (y_S)_{S \in \mathcal{V}_{n,2d}}$, and its associated pseudo-moment matrix is $H_d(y)$.

We define a linear subspace of matrices with rows and columns indexed by $\mathcal{V}_{n,d}$ by imposing two kinds of linear constraints.

(1) *Equal-union constraints:*

$$A_{S,T} = A_{S',T'} \qquad \text{whenever } S \cup T = S' \cup T'. \tag{9}$$

After imposing these constraints, for every $R \in \mathcal{V}_{n,2d}$ we write y_R for the common value $A_{S,T}$ over all pairs $S, T \in \mathcal{V}_{n,d}$ with $S \cup T = R$. For convenience, we describe the next set of constraints in terms of y.

(2) *Localizing constraints:* For each QuadEq equation

$$f_\ell(x) = \sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} x^U, \qquad c_{\ell,U} \in \mathbb{F}_q,$$

we impose the linear constraints

$$\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U}\, y_{U \cup W} = 0 \qquad \text{for every } \ell \in [m] \text{ and every } W \in \mathcal{V}_{n,2d-2}. \tag{10}$$

These are the moment analogues of the equations $x^W f_\ell(x) = 0$: after multiplying f_ℓ by any squarefree monomial of degree at most $2d-2$, the corresponding linear combination of pseudo-moments must vanish.

Let $\mathcal{L}_d(f_1, \ldots, f_m)$ denote the set of matrices satisfying the constraints above. This is a linear subspace of $\mathbb{F}_q^{N \times N}$, where

$$N = |\mathcal{V}_{n,d}| = \sum_{j=0}^{d} \binom{n}{j} = n^{O(d)}.$$

For completeness, if $a = (a_1, \ldots, a_n) \in \{0,1\}^n$ satisfies all equations $f_\ell(a) = 0$, set $y_S := a^S$ for every $S \in \mathcal{V}_{n,2d}$. Then $H_d(y) = v_d(a)v_d(a)^\top$ has rank 1 and is nonzero because the coordinate indexed by $\emptyset$ is 1. Moreover, for every $W \in \mathcal{V}_{n,2d-2}$,

$$\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U}\, y_{U \cup W} = a^W f_\ell(a) = 0,$$

so $H_d(y) \in \mathcal{L}_d(f_1, \ldots, f_m)$.

We now begin the soundness proof. It suffices to show that every nonzero matrix in $\mathcal{L}_d(f_1, \ldots, f_m)$ with rank at most d yields an actual Boolean solution of the source instance. Let y be the unique pseudo-moment vector associated with such a matrix A, so $A = H_d(y)$ and y satisfies (10). For a genuine Boolean point a, multiplying the label of a monomial column by x_i multiplies that column by the scalar a_i; equivalently, the column space is invariant under multiplication by each variable. We will find two consecutive levels where low rank forces a similar form of invariance.

For $0 \le e \le d$, let

$$H_e(y) := \big(y_{S \cup T}\big)_{S,T \in \mathcal{V}_{n,e}}, \qquad r_e := \operatorname{rank} H_e(y),$$

and let C_e denote the column space of $H_e(y)$.

The soundness argument has four steps.

(1) Low rank forces a nonzero *flat level*, meaning a level e with $r_e = r_{e+1} > 0$.

(2) At such a flat level, multiplying a basis column by any variable x_i creates no new column direction. This defines a linear map $T_i : C_e \to C_e$, interpreted as multiplication by x_i.

(3) Because the relevant column relations are certified one level up (a.k.a inside C_{e+1}), these maps satisfy the Boolean multiplication rules $T_i^2 = T_i$ and $T_iT_j = T_jT_i$, and the localizing equations become the operator identities $f_\ell(T) = 0$.

(4) The commuting projections T_i have a common eigenvector. Their eigenvalues form a Boolean point $a \in \{0,1\}^n \subseteq \mathbb{F}_q^n$, and the identities $f_\ell(T) = 0$ force $f_\ell(a) = 0$.

4.2 Finding a nonzero flat level

Note that the ranks $r_0, r_1, \ldots, r_d$ are non-decreasing, and the final rank is at most d. Monotonicity alone would still allow the ranks to be zero for many initial levels and then grow by one at each remaining level. The key point is that a nonzero pseudo-moment cannot start growing this slowly. If S is a minimum-size set with $y_S \neq 0$, then the same permutation-submatrix idea as in Lemma 3.10 already forces, at level $s = \lceil |S|/2 \rceil$, the lower bound $r_s \ge s+1$ (indeed $r_s \ge 2^{\Omega(s)}$ except for the smallest cases). If every subsequent positive step were strict, then

$$r_d \ge r_s + (d-s) \ge (s+1) + (d-s) = d+1,$$

contradicting $r_d \le d$. Thus some nonzero plateau $r_e = r_{e+1}$ must occur.

▶ **Lemma 4.3.** *Let $d \geq 1$. Suppose $H_d(y)$ is nonzero and* $\operatorname{rank} H_d(y) \leq d$. *Then there exists $e \in \{0, 1, \dots, d-1\}$ such that*

$$r_e = r_{e+1} > 0.$$

Proof. Choose a set $S \in \mathcal{V}_{n,2d}$ of minimum size such that $y_S \neq 0$, and put

$$\lambda := y_S, \qquad h := |S|, \qquad s := \lceil h/2 \rceil .$$

Since $h \leq 2d$, we have $s \leq d$. The first step is the following rank lower bound forced by the first nonzero moment:

$$r_s \geq s + 1. \tag{11}$$

If $h = 0$, then $s = 0$ and $y_\emptyset = \lambda \neq 0$, so $r_s = r_0 = 1 = s + 1$. If $h = 1$, say $S = \{i\}$, then the minimality of S gives $y_\emptyset = 0$ and $y_S = \lambda$. The submatrix of $H_1(y)$ with rows and columns indexed by $\emptyset, S$ is

$$\begin{pmatrix} y_\emptyset & y_S \\ y_S & y_S \end{pmatrix} = \begin{pmatrix} 0 & \lambda \\ \lambda & \lambda \end{pmatrix},$$

which has rank $2 = s + 1$.

It remains to consider $h \geq 2$. Use the subsets $A \subseteq S$ of size $\lfloor h/2 \rfloor$ as row labels and the subsets $B \subseteq S$ of size $\lceil h/2 \rceil = s$ as column labels. The resulting submatrix of $H_s(y)$ has entries $y_{A \cup B}$. By the minimality of S, this entry is λ exactly when $A \cup B = S$, and is 0 otherwise. For each row A, there is a unique such column, namely $B = S \setminus A$. Thus this submatrix is λ times a permutation matrix of size

$$\binom{h}{\lfloor h/2 \rfloor}.$$

Therefore

$$r_s \geq \binom{h}{\lfloor h/2 \rfloor} \geq h \geq \lceil h/2 \rceil + 1 = s + 1,$$

where the middle inequality uses $h \geq 2$. This proves (11) in all cases.

The matrices $H_0(y), H_1(y), \dots, H_d(y)$ are nested principal submatrices, so the ranks $r_0, r_1, \dots, r_d$ are nondecreasing. Suppose, for contradiction, that no $e \in \{0, 1, \dots, d-1\}$ satisfies $r_e = r_{e+1} > 0$. Since $r_s > 0$ and the ranks are nondecreasing, the ranks strictly increase at every such step. Thus

$$r_d \geq r_s + (d - s) \geq (s + 1) + (d - s) = d + 1,$$

contradicting $r_d = \operatorname{rank} H_d(y) \leq d$. Therefore a nonzero flat level exists. ◀

4.3 Multiplication operators and localizing identities

For $0 \leq e \leq d$ and a set $A \subseteq [n]$ with $|A| \leq 2d - e$, define the truncated column

$$\mathbf{c}_e(A) := \big(y_{R \cup A}\big)_{R \in \mathcal{V}_{n,e}} \in \mathbb{F}_q^{\mathcal{V}_{n,e}}.$$

If $A \in \mathcal{V}_{n,e}$, then $\mathbf{c}_e(A)$ is exactly the column of $H_e(y)$ indexed by A. If $|A| = t > e$, it is not literally a column of $H_e(y)$, but it is still the degree-e truncation of the column indexed by A, in a larger moment matrix $H_t(y)$.

We now explain how the flat level is used. At a flat level $r_e = r_{e+1}$, any basis of the column space C_e remains a basis after lifting the columns to level $e+1$. Hence multiplying the column label by x_i gives a well-defined operator $T_i : C_e \to C_e$. Furthermore, this extra-level lift certifies that

$$T_i \mathbf{c}_e(A) = \mathbf{c}_e(A \cup \{i\})$$

for every $A \in \mathcal{V}_{n,e+1}$. This allows us to multiply the column label by a second variable consistently, yielding $T_i^2 = T_i$, $T_iT_j = T_jT_i$, and ultimately translating the localizing constraints into the operator identities $f_\ell(T) = 0$.

▶ **Lemma 4.4.** *Assume $r_e = r_{e+1}$ and $e \leq d-1$. Let $\mathcal{I} \subseteq \mathcal{V}_{n,e}$ be such that the vectors $\{\mathbf{c}_e(B) : B \in \mathcal{I}\}$ form a basis for C_e. Then, for every $i \in [n]$, there is a linear map $T_i : C_e \to C_e$ satisfying*

$$T_i \mathbf{c}_e(B) = \mathbf{c}_e(B \cup \{i\}) \qquad \text{for every } B \in \mathcal{I}.$$

Moreover, for every $A \in \mathcal{V}_{n,e+1}$ and every $i \in [n]$, the vector $\mathbf{c}_e(A)$ lies in C_e and

$$T_i \mathbf{c}_e(A) = \mathbf{c}_e(A \cup \{i\}). \tag{12}$$

Proof. First observe that the longer vectors $\{\mathbf{c}_{e+1}(B) : B \in \mathcal{I}\}$ are linearly independent: any linear relation among them restricts, on the rows indexed by $\mathcal{V}_{n,e}$, to the same linear relation among the basis vectors $\{\mathbf{c}_e(B) : B \in \mathcal{I}\}$. Since $r_e = r_{e+1}$, these longer vectors form a basis for the column space of $H_{e+1}(y)$.

For every $i \in [n]$ and every $B \in \mathcal{I}$, the vector $\mathbf{c}_{e+1}(B \cup \{i\})$ is a column of $H_{e+1}(y)$, and hence has a unique expansion in the basis $\{\mathbf{c}_{e+1}(P) : P \in \mathcal{I}\}$. Restricting that expansion to rows of degree at most e shows that $\mathbf{c}_e(B \cup \{i\}) \in C_e$. We may therefore define T_i on the basis by

$$T_i \mathbf{c}_e(B) := \mathbf{c}_e(B \cup \{i\}) \qquad (B \in \mathcal{I}),$$

and extend linearly to all of C_e.

Now fix $A \in \mathcal{V}_{n,e+1}$. We would like to argue that $T_i \mathbf{c}_e(A) = \mathbf{c}_e(A \cup \{i\})$. Since $\mathbf{c}_{e+1}(A)$ is a column of $H_{e+1}(y)$, there are unique coefficients $\lambda_{A,B} \in \mathbb{F}_q$ such that

$$\mathbf{c}_{e+1}(A) = \sum_{B \in \mathcal{I}} \lambda_{A,B} \mathbf{c}_{e+1}(B).$$

Restricting to rows in $\mathcal{V}_{n,e}$ gives

$$\mathbf{c}_e(A) = \sum_{B \in \mathcal{I}} \lambda_{A,B} \mathbf{c}_e(B),$$

so $\mathbf{c}_e(A) \in C_e$. By linearity and the definition of T_i on the basis,

$$T_i \mathbf{c}_e(A) = \sum_{B \in \mathcal{I}} \lambda_{A,B} T_i \mathbf{c}_e(B) = \sum_{B \in \mathcal{I}} \lambda_{A,B} \mathbf{c}_e(B \cup \{i\}).$$

It remains to identify the last expression with $\mathbf{c}_e(A \cup \{i\})$. Fix a row label $R \in \mathcal{V}_{n,e}$. Since $R \cup \{i\} \in \mathcal{V}_{n,e+1}$, evaluating the expansion of $\mathbf{c}_{e+1}(\cdot)$ at the row $R \cup \{i\}$ gives

$$\begin{aligned}
\sum_{B \in \mathcal{I}} \lambda_{A,B} \big(\mathbf{c}_e(B \cup \{i\})\big)_R &= \sum_{B \in \mathcal{I}} \lambda_{A,B} y_{R \cup B \cup \{i\}} \\
&= \sum_{B \in \mathcal{I}} \lambda_{A,B} \big(\mathbf{c}_{e+1}(B)\big)_{R \cup \{i\}} \\
&= \big(\mathbf{c}_{e+1}(A)\big)_{R \cup \{i\}} \\
&= y_{R \cup A \cup \{i\}} = \big(\mathbf{c}_e(A \cup \{i\})\big)_R .
\end{aligned}$$

This holds for every row $R \in \mathcal{V}_{n,e}$, so $T_i \mathbf{c}_e(A) = \mathbf{c}_e(A \cup \{i\})$. The vector on the right is well-defined because $|A \cup \{i\}| \leq e + 2 \leq 2d - e$, using $e \leq d - 1$. ◀

▶ **Lemma 4.5.** *Assume $r_e = r_{e+1} > 0$ and $e \le d-1$. Let C_e be the column space of $H_e(y)$. If y satisfies the equations* (10)*, then there are linear maps*

$$T_1, \dots, T_n : C_e \to C_e$$

with the following properties:

(i) $T_i^2 = T_i$ *for every* i*;*

(ii) $T_i T_j = T_j T_i$ *for every* i, j*;*

(iii) *for every* $\ell \in [m]$*,*

$$\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} T_U = 0 \qquad \text{as a linear map on } C_e,$$

where $T_U := \prod_{i \in U} T_i$ *and* $T_\emptyset$ *is the identity map.*

Proof. Choose $\mathcal{I} \subseteq \mathcal{V}_{n,e}$ such that $\{\mathbf{c}_e(B) : B \in \mathcal{I}\}$ is a basis for C_e, and define the maps T_i as in Lemma 4.4. By Lemma 4.4, for every $A \in \mathcal{V}_{n,e+1}$ and every i,

$$T_i \mathbf{c}_e(A) = \mathbf{c}_e(A \cup \{i\}). \tag{13}$$

We first prove the Boolean multiplication rules. Fix $B \in \mathcal{I}$ and $i, j \in [n]$. Taking $A = B$ in (13) gives $T_i \mathbf{c}_e(B) = \mathbf{c}_e(B \cup \{i\})$. Since $B \cup \{i\} \in \mathcal{V}_{n,e+1}$, applying (13) with $A = B \cup \{i\}$ and variable j gives

$$T_j T_i \mathbf{c}_e(B) = T_j \mathbf{c}_e(B \cup \{i\}) = \mathbf{c}_e(B \cup \{i, j\}). \tag{14}$$

Taking $j = i$ in (14) gives $T_i^2 \mathbf{c}_e(B) = T_i \mathbf{c}_e(B)$, because $B \cup \{i, i\} = B \cup \{i\}$. Since this holds on the basis vectors $\{\mathbf{c}_e(B) : B \in \mathcal{I}\}$, we get $T_i^2 = T_i$. Swapping i and j in (14) gives

$$T_i T_j \mathbf{c}_e(B) = \mathbf{c}_e(B \cup \{j, i\}) = \mathbf{c}_e(B \cup \{i, j\}) = T_j T_i \mathbf{c}_e(B),$$

again on every basis vector. Hence $T_i T_j = T_j T_i$.

It remains to translate the original equations into operator identities. Because the T_i commute, $T_U := \prod_{i \in U} T_i$ is well-defined. From (13) when $|U| \le 1$, and from (14) when $|U| = 2$, we have

$$T_U \mathbf{c}_e(B) = \mathbf{c}_e(B \cup U) \qquad \text{for every } B \in \mathcal{I} \text{ and every } U \in \mathcal{V}_{n,2}.$$

Therefore, for every row label $R \in \mathcal{V}_{n,e}$,

$$\left(\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} T_U \mathbf{c}_e(B) \right)_R = \sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} y_{R \cup B \cup U}.$$

Here $R \cup B$ has size at most $2e$. Since $e \le d-1$, we have $2e \le 2d-2$, so the localizing equations (10) apply with $W = R \cup B$. Hence the last sum is 0. Thus $\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} T_U$ kills every basis vector of C_e, and therefore it is the zero map on C_e. ◀

4.4 Decoding a satisfying assignment from commuting projections

The next lemma records the common-eigenvector property for commuting projections. We omit the proof here.

▶ **Lemma 4.6.** *Let C be a nonzero finite-dimensional vector space over $\mathbb{F}_q$, and let $T_1, \ldots, T_n$ be commuting linear maps on C with $T_i^2 = T_i$ for every i. Then there are $v \in C \setminus \{0\}$ and scalars $a_i \in \{0,1\} \subseteq \mathbb{F}_q$, $i \in [n]$, such that*

$$T_i v = a_i v \qquad \text{for every } i = 1, \ldots, n.$$

▶ **Lemma 4.7.** *Let C be a nonzero finite-dimensional vector space over $\mathbb{F}_q$, and let $T_1, \ldots, T_n$ be commuting linear maps on C. Suppose that, for every $\ell \in [m]$,*

$$\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} T_U = 0 \qquad \text{as a linear map on } C,$$

where $T_U := \prod_{i \in U} T_i$ and $T_\emptyset$ is the identity map. If there are $v \in C \setminus \{0\}$ and scalars $a_1, \ldots, a_n \in \{0,1\} \subseteq \mathbb{F}_q$ such that $T_i v = a_i v$ for every i, then $a = (a_1, \ldots, a_n)$ satisfies $f_1(a) = \cdots = f_m(a) = 0$.

Proof. Since the operators commute, for every $U \in \mathcal{V}_{n,2}$ we have

$$T_U v = a^U v.$$

Therefore, for every ℓ,

$$0 = \Big(\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} T_U \Big) v = \Big(\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} a^U \Big) v = f_\ell(a) v.$$

Since $v \neq 0$, the scalar $f_\ell(a)$ must be zero. Hence the Boolean point $a = (a_1, \ldots, a_n)$ satisfies all equations. ◀

We are now ready to prove Theorem 4.2.

Proof of Theorem 4.2. Given the Boolean QuadEq instance $f_1, \ldots, f_m$, set $d := k$ and output

$$\mathcal{L} := \mathcal{L}_d(f_1, \ldots, f_m).$$

The matrix size is

$$N = |\mathcal{V}_{n,d}| = \sum_{j=0}^{k} \binom{n}{j} = n^{O(k)},$$

and the number of equal-union and localizing constraints is also $(n+m)n^{O(k)}$, so the construction has the claimed running time.

If the Boolean QuadEq instance has a solution, the completeness argument above shows that $\mathcal{L}$ contains a nonzero rank-1 matrix. If the Boolean QuadEq instance has no solution and $A \in \mathcal{L}$ is nonzero with $\operatorname{rank}(A) \le k$, then the unique pseudo-moment vector y associated with A satisfies $A = H_k(y)$ and (10). By Lemma 4.3, choose a flat level e with $r_e = r_{e+1} > 0$. By Lemma 4.5, the column space C_e carries commuting projections $T_1, \ldots, T_n$, and for every $\ell \in [m]$,

$$\sum_{U \in \mathcal{V}_{n,2}} c_{\ell,U} T_U = 0.$$

By Lemma 4.6, there are a nonzero vector $v \in C_e$ and scalars $a_1, \ldots, a_n \in \{0,1\} \subseteq \mathbb{F}_q$ such that $T_i v = a_i v$ for every i. By Lemma 4.7, a is a Boolean solution, a contradiction. Hence no such nonzero matrix exists. ◀

The inapproximability statement Theorem 1.2 is a standard corollary of Theorem 4.2.

References

1 Per Austrin and Subhash Khot. A simple deterministic reduction for the gap minimum distance of code problem. *IEEE Transactions on Information Theory*, 60(10):6636–6645, 2014. doi:10.1109/TIT.2014.2340869.

2 Boaz Barak and David Steurer. Sum-of-squares proofs and the quest toward optimal algorithms. In *Proceedings of the International Congress of Mathematicians (ICM)*, 2014. arXiv:1404.5236.

3 Vijay Bhattiprolu, Venkatesan Guruswami, Euiwoong Lee, and Xuandi Ren. Inapproximability of finding sparse vectors in codes, subspaces, and lattices. In *66th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2025, Sydney, Australia, December 14-17, 2025*, pages 1295–1303. IEEE, 2025. doi:10.1109/FOCS63196.2025.00068.

4 Qi Cheng and Daqing Wan. A deterministic reduction for the gap minimum distance problem. *IEEE Transactions on Information Theory*, 58(11):6935–6941, 2012. doi:10.1109/TIT.2012.2209198.

5 Nicolas T. Courtois. Efficient zero-knowledge authentication based on a linear algebra problem MinRank. In *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 402–421. Springer, 2001. doi:10.1007/3-540-45682-1_24.

6 Irit Dinur and Venkatesan Guruswami. Pcps via the low-degree long code and hardness for constrained hypergraph coloring. *Israel Journal of Mathematics*, 209(2):611–649, 2015. doi:10.1007/s11856-015-1231-3.

7 Ilya Dumer, Daniele Micciancio, and Madhu Sudan. Hardness of approximating the minimum distance of a linear code. *IEEE Transactions on Information Theory*, 49(1):22–37, 2003. doi:10.1109/TIT.2002.806118.

8 Aviezri S. Fraenkel and Yaacov Yesha. Complexity of problems in games, graphs and algebraic equations. *Discrete Applied Mathematics*, 1(1–2):15–30, 1979. doi:10.1016/0166-218X(79)90012-X.

9 Philippe Gaborit and Gilles Zémor. On the hardness of the decoding and the minimum distance problems for rank codes. *IEEE Transactions on Information Theory*, 62(12):7245–7252, 2016. doi:10.1109/TIT.2016.2616127.

10 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness.* W. H. Freeman, 1979.

11 Sangxia Huang. $2^{(\log N)^{1/10-o(1)}}$ hardness for hypergraph coloring. *arXiv preprint arXiv:1504.03923*, 2015. doi:10.48550/arXiv.1504.03923.

12 Subhash Khot and Rishi Saket. Hardness of coloring 2-colorable 12-uniform hypergraphs with $2^{(\log n)^{\Omega(1)}}$ colors. In *55th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2014)*, pages 206–215, 2014. doi:10.1109/FOCS.2014.30.

13 Jean B. Lasserre. Global optimization with polynomials and the problem of moments. *SIAM Journal on Optimization*, 11(3):796–817, 2001. doi:10.1137/S1052623400366802.

14 Monique Laurent. Sums of squares, moment matrices and optimization over polynomials. In Mihai Putinar and Seth Sullivant, editors, *Emerging Applications of Algebraic Geometry*, volume 149 of *The IMA Volumes in Mathematics and its Applications*, pages 157–270. Springer, 2009. doi:10.1007/978-0-387-09686-5_7.

15 Daniele Micciancio. Locally dense codes. In *2014 IEEE 29th Conference on Computational Complexity (CCC)*, pages 90–97. IEEE, 2014. doi:10.1109/CCC.2014.17.

16 Pablo A. Parrilo. Semidefinite programming relaxations for semialgebraic problems. *Mathematical Programming*, 96:293–320, 2003. doi:10.1007/s10107-003-0387-5.

17 Alberto Ravagnani. Rank-metric codes and their duality theory. *Designs, Codes and Cryptography*, 80(1):197–216, 2016. doi:10.1007/s10623-015-0077-3.