

Deterministic Hardness-of-Approximation of Unique-SVP and GapSVP in ℓ_p norms for $p > 2$

Yahli Hecht*

Muli Safra†

2025

Abstract

We establish *deterministic* hardness of approximation results for the Shortest Vector Problem in ℓ_p norm (SVP_p) and for Unique-SVP (uSVP_p)—namely, instances promised to have a unique shortest vector—for all $p > 2$. Previously, no *deterministic* hardness results were known, except for ℓ_∞ .

For every $p > 2$, we prove constant-ratio hardness: no polynomial-time algorithm approximates SVP_p or uSVP_p within a ratio of $\sqrt{2} - o(1)$, assuming $3\text{SAT} \notin \text{DTIME}(2^{O(n^{2/3} \log n)})$, and, respectively, $\text{Unambiguous-3SAT} \notin \text{DTIME}(2^{O(n^{2/3} \log n)})$.

We also show that for any $\varepsilon > 0$ there exists $p_\varepsilon > 2$ such that for every $p \geq p_\varepsilon$: no polynomial-time algorithm approximates SVP_p within a ratio of $2^{(\log n)^{1-\varepsilon}}$, assuming $\text{NP} \not\subseteq \text{DTIME}(n^{(\log n)^\varepsilon})$; and within a ratio of $n^{1/(\log \log n)^\varepsilon}$, assuming $\text{NP} \not\subseteq \text{SUBEXP}$. This improves upon [Haviv, Regev, Theory of Computing 2012], which obtained similar inapproximation ratios under randomized reductions. We obtain analogous results for uSVP_p under the assumptions $\text{Unambiguous-3SAT} \not\subseteq \text{DTIME}(n^{(\log n)^\varepsilon})$ and $\text{Unambiguous-3SAT} \not\subseteq \text{SUBEXP}$, improving the previously known $1 + o(1)$ [Stephens-Davidowitz, Approx 2016].

Strengthening the hardness of uSVP at weaker approximation ratios has direct cryptographic impact. By the reduction of Lyubashevsky and Micciancio [Lyubashevsky, Micciancio, CRYPTO 2009], hardness for $\gamma\text{-uSVP}_p$ carries over to $\frac{1}{\gamma}\text{-BDD}_p$ (Bounded Distance Decoding). Thus, understanding the hardness of uSVP improves worst-case guarantees for the two core problems that underpin security in lattice-based cryptography.

1 Introduction

A lattice $\mathcal{L} \subseteq \mathbb{R}^n$ is the additive group of all integer linear combinations of d linearly independent vectors. Given a linearly independent matrix $M \in \mathbb{R}^{n \times d}$, we write $\mathcal{L}[M] \stackrel{\text{def}}{=} \{M \cdot \vec{a} \mid \vec{a} \in \mathbb{Z}^d\}$. For $p \in [1, \infty]$, the ℓ_p norm is $\|x\|_p = (\sum_i |x_i|^p)^{1/p}$ for $p < \infty$ and $\|x\|_\infty = \max_i |x_i|$. The length of the shortest nonzero vector is $\lambda_1^{(p)}(\mathcal{L}) \stackrel{\text{def}}{=} \min_{v \in \mathcal{L} \setminus \{0\}} \|v\|_p$, equivalently, the smallest r such that the closed ℓ_p -ball $B_p(0, r)$ contains a nonzero lattice point. More generally, the k -th successive minimum $\lambda_k^{(p)}(\mathcal{L})$ is the least r for which $B_p(0, r)$ contains k linearly independent lattice vectors.

The main goal of this paper is to establish deterministic hardness of approximation results for two lattice problems, known as SVP and uSVP . In the $\gamma\text{-SVP}_p$ problem ($\gamma > 1$), we are given a basis M for $\mathcal{L}[M]$ and a

*School of Computer Science, Tel Aviv University. Supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (Grant agreement No. 835152). email: yahlihecht@mail.tau.ac.il

†School of Computer Science, Tel Aviv University. Supported by the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation programme (Grant agreement No. 835152) and by the Israel Science Foundation (ISF) grant 2257/21. email: safra@mail.tau.ac.il

radius r , and the goal is to distinguish between the case of $\lambda_1^{(p)}(\mathcal{L}[M]) \leq r$ and the case of $\lambda_1^{(p)}(\mathcal{L}[M]) \geq \gamma r$. In the γ -uSVP $_p$ problem, the YES instances are promised to satisfy $\lambda_2^{(p)}(\mathcal{L}[M]) \geq \gamma \lambda_1^{(p)}(\mathcal{L}[M])$. In words, the shortest vector in $\mathcal{L}[M]$ is unique—the only non-zero vectors in the lattice that are of length less than $\gamma \lambda_1(\mathcal{L}[M])$ are its multiples. A closely related problem is BDD. In the $\frac{1}{\gamma}$ -BDD $_p$, we are given a basis M and a target vector $\vec{t}$ such that $\text{dist}(\mathcal{L}[M], \vec{t}) \leq \frac{1}{\gamma} \lambda_1^{(p)}$, find the lattice vector closest to $\vec{t}$. Formal definitions appear in the preliminaries (Section 2).

Hardness of SVP. The study of the computational hardness of the shortest vector problem (SVP) has a long history. The first hardness result, due to van Emde Boas [vE81], established NP-hardness of SVP in the ℓ_∞ norm. Ajtai [Ajt98] proved NP-hardness in ℓ_2 , for an inapproximation ratio slightly larger than 1 and via a *randomized reduction*. Micciancio [Mic98, Mic01] improved the inapproximability ratio within $2^{1/p} - o(1)$ for all $1 \leq p < \infty$. For high p , Khot [Kho03] proved hardness of approximation to within $p^{1-\epsilon}$. Khot [Kho05] later achieved the hardness for every constant inapproximation ratio, by constructing GapSVP instances with an additional structure and utilizing a variant of the tensor product to amplify the gap. Both Micciancio’s and Khot’s reductions are *randomized*.

Allowing random quasi-polynomial reductions, Khot [Kho05] also established hardness for a ratio of $2^{\log^{\frac{1}{2}-\epsilon} n}$. Haviv and Regev [HR12, HR18] improved the ratio to $2^{\log^{1-\epsilon} n}$. Under the stronger yet plausible assumption $\text{NP} \not\subseteq \text{RSUBEXP} = \cap_{\gamma > 0} \text{RTIME}(2^{n^\gamma})$ the ratio reaches $n^{O(1/\log \log n)}$.

Hardness of Unique-SVP. For the *unique* variant, Kumar and Sivakumar [KS01] first proved NP-hardness in ℓ_2 , albeit under randomized reductions. Khoat and Tan [KT08] proved NP-hardness of exact uSVP in ℓ_∞ . In ℓ_p , the best unconditional hardness factors remain very close to one: Aggarwal and Dubey [AD16] proved hardness within $1 + 1/\text{poly}(n)$, and Stephens-Davidowitz [SD15] achieved $1 + O(\log n/n)$. More recently, Jin and Xue [JX24] presented a fine-grained hardness result, reaching an inapproximability ratio of $1 + \epsilon$. Under certain nonstandard assumptions, [BGPSD23] established hardness for every constant ratio $\gamma \geq 1$. Lyubashevsky and Micciancio [LM09] provided further evidence for the hardness of uSVP, by proving equivalence to GapSVP up to a small polynomial factor of $\sqrt{n/\log n}$.

1.1 Our results

All of the above mentioned hardness of approximation results for SVP $_p$ and uSVP $_p$ have relied on *randomized* reductions (beyond ℓ_∞). At the cost of restricting attention to $p > 2$, we obtain deterministic reductions that match—and sometimes substantially improve upon—existing hardness results.

For SVP, we prove hardness of approximation within $\sqrt{2} - o(1)$ for every $p > 2$, under deterministic sub-exponential reductions. When p is sufficiently large, we deterministically match and strengthen the best known randomized results in the high- p regime, improving the $n^{O(1/\log \log n)}$ ratio of Haviv and Regev [HR18].

Theorem 1.1 (SVP, $p > 2$). *For every constant $p > 2$, deciding GapSVP in ℓ_p is hard to approximate within a ratio $\sqrt{2} - o(1)$, unless $3\text{SAT} \in \text{DTIME}\left(2^{O(n^{2/3} \log n)}\right)$.*

Theorem 1.2 (SVP, high p). *For every $\epsilon > 0$ there exists $p_\epsilon > 2$ such that for every $p \geq p_\epsilon$, GapSVP in ℓ_p is hard to approximate within a ratio $2^{(\log n)^{1-\epsilon}}$, unless $\text{NP} \subseteq \text{DTIME}(n^{(\log n)^\epsilon})$. Furthermore, under the stronger assumption $\text{NP} \not\subseteq \text{SUBEXP}$, GapSVP is hard to approximate within a ratio $n^{1/(\log \log n)^\epsilon}$.*

Prior hardness results for uSVP lag far behind those for SVP, failing to reach even constant inapproximability ratios. Beyond ℓ_∞ , existing reductions are again *randomized*. Our reductions *substantially* improve this picture. For every $p > 2$, we give a *deterministic* reduction within a ratio of $(\sqrt{2} - o(1))$. For sufficiently large p , we obtain a quasi-polynomial reduction showing hardness of approximation within an almost-polynomial inapproximability factor. Thus, in the uSVP regime we move from sub-constant factors (best at $1 + \frac{\log n}{n}$ by Stephens-Davidowitz [SD15]) to almost-polynomial factors aligning with the known picture for SVP.

Result	ℓ_p	Approx. ratio	Assumption	Prior works
Theorem 1.1	$p > 2$	$\sqrt{2} - o(1)$	$3\text{SAT} \notin \text{DTIME}\left(2^{O(n^{2/3} \log n)}\right)$	—
Theorem 1.2	$p \geq p_\varepsilon$	$2^{(\log n)^{1-\varepsilon}}$	$\text{NP} \not\subseteq \text{DTIME}\left(n^{(\log n)^\varepsilon}\right)$	$2^{(\log n)^{1-\varepsilon}}$; $\text{NP} \not\subseteq \text{RTIME}\left(n^{(\log n)^c}\right)$; [HR18]
Theorem 1.2	$p \geq p_\varepsilon$	$n^{1/(\log \log n)^\varepsilon}$	$\text{NP} \not\subseteq \text{SUBEXP}$	$n^{O(1/(\log \log n))}$; $\text{NP} \not\subseteq \text{RSUBEXP}$; [HR18]

For uSVP , our results are reduced from Unambiguous-3SAT . In Unambiguous-3SAT , the task is to distinguish 3SAT formulas that have exactly one satisfying assignment from those that are unsatisfiable. By the Valiant–Vazirani theorem [\[VV85\]](#), no polynomial-time algorithm decides Unambiguous-3SAT unless $\text{NP} \subseteq \text{RP}$.

Theorem 1.3 (uSVP , $p > 2$). *For every constant $p > 2$, uSVP in ℓ_p is hard to approximate within a ratio $\sqrt{2} - o(1)$, unless $\text{Unambiguous-3SAT} \in \text{DTIME}\left(2^{O(n^{2/3} \log n)}\right)$.*

Theorem 1.4 (uSVP , high p). *For any $\varepsilon > 0$ there exists $p_\varepsilon > 2$ so that for every $p \geq p_\varepsilon$, uSVP in ℓ_p is hard to approximate within a ratio $2^{(\log n)^{1-\varepsilon}}$, unless $\text{Unambiguous-3SAT} \subseteq \text{DTIME}\left(n^{(\log n)^\varepsilon}\right)$. Under the stronger assumption $\text{Unambiguous-3SAT} \not\subseteq \text{SUBEXP}$, uSVP is hard to approximate within a ratio $n^{1/(\log \log n)^\varepsilon}$.*

Result	p -range	Approx. ratio	Assumption	Prior works
Theorem 1.3	$p > 2$	$\sqrt{2} - o(1)$	$\text{Unambiguous-3SAT} \notin \text{DTIME}\left(2^{O(n^{2/3} \log n)}\right)$	$1 + \delta$; running-time $2^{\varepsilon n}$ (from SVP); [JX24]
Theorem 1.4	$p \geq p_\varepsilon$	$2^{(\log n)^{1-\varepsilon}}$	$\text{Unambiguous-3SAT} \not\subseteq \text{DTIME}\left(n^{(\log n)^\varepsilon}\right)$	$1 + \frac{\log n}{n}$; $\text{NP} \not\subseteq \text{RP}$; [SD15]
Theorem 1.4	$p \geq p_\varepsilon$	$n^{1/(\log \log n)^\varepsilon}$	$\text{Unambiguous-3SAT} \not\subseteq \text{SUBEXP}$	...

1.2 Motivation

Determinism. Beyond the ℓ_∞ norm, essentially all known hardness-of-approximation results for SVP and uSVP rely on *randomized* reductions, leaving open whether randomness is inherently necessary. Derandomizing these reductions has therefore been a long-standing goal. Previous efforts have built upon the randomized reductions to SVP_2 , introducing alternative gadget constructions that may be easier to derandomize. In this line, Micciancio [\[Mic12\]](#) obtained a reduction achieving the inapproximability ratio of [\[HR18\]](#) with one-sided error, and Bennett and Peikert [\[BP23\]](#) explored deterministic gadgets based on Reed–Solomon codes.

We take a different route: returning to the algebraic PCP framework of [\[DFK⁺99, Din02, DKRS03\]](#), and adapting its lattice encodings of NP witnesses as short lattice vectors. We generalize this machinery and make it applicable to SVP in ℓ_p for $p > 2$. This framework is modular and we expect it to yield additional derandomized hardness results for lattice problems, possibly extending even to the ℓ_2 case.

Cryptography and hardness of approximation. Tightening the approximation ratio for uSVP has direct consequences for both worst-case hardness and cryptography. Lyubashevsky and Micciancio [\[LM09\]](#) showed that, for any $p \geq 1$,

$$\gamma\text{-uSVP}_p \leq \frac{1}{\gamma}\text{-BDD}_p \leq 2\gamma\text{-uSVP}_p,$$

so the security of lattice-based cryptosystems whose assumptions reduce to either BDD or uSVP directly depends on their respective approximation hardness. Examples include encryption and signature schemes

based on **LWE** [Reg09, Pei16], **NTRU** [HPS98], and **SIS/Ajtai-Dwork-type** constructions [AD97, GGH96, Reg04, AD07]. The complexity of approximating **BDD** is somewhat better understood than that of **uSVP**. Liu, Lyubashevsky and Micciancio [LLM06] established NP-hardness of BDD_p for small constant approximation factors (for any $p \geq 1$), and Bennett and Peikert [BP20] showed NP-hardness for $\alpha\text{-BDD}_p$, with ratios $\alpha \rightarrow \frac{1}{2}$ as $p \rightarrow \infty$, approaching the unique-decoding radius (at most a single close vector exists—as lattice vectors are at distance of at least $\lambda_1^{(p)}$). Surpassing this barrier provides additional motivation.

Attacks against LWE. **BDD** is precisely the decoding task underlying **LWE** (Learning With Errors) [Reg09]. A standard attack pipeline treats **LWE** as an instance of **BDD** and reduces it to **uSVP**, commonly via Kannan’s embedding technique [Kan87]:

$$\text{LWE} \longrightarrow \text{BDD} \longrightarrow \text{uSVP}.$$

This pathway and connected attacks have been analyzed and optimized in many works, including [LP11, CN11, AFG13, BSW16, AGVW17]. Cryptographic constructions assume hardness of approximation ratios for **uSVP** that are much larger than the regime where NP-hardness is known or believed (see [GG98, Cai98, AR05], for strong evidence against it). Nevertheless, this motivates sharper hardness of approximation results for both **BDD** and **uSVP**.

1.3 Outline

The paper is organized as follows. [Section 2](#) reviews standard definitions and tools. [Section 3](#) presents core components of our constructions, combining prior work with our modifications and gadgets. The sub-exponential construction for [Theorem 1.1](#) and [Theorem 1.3](#) appears in [Section 4](#), together with its soundness analysis. In [Section 5](#), we give our construction for [Theorem 1.2](#) and [Theorem 1.4](#). [Section 6](#) presents its soundness analysis. In [Section 7](#), we conclude with a discussion and present some open problems for future research.

2 Preliminaries

In this section, we recall standard definitions and tools for lattices and lattice problems. We assume familiarity with basic concepts from PCPs, lattice geometry, and probability. For a thorough exposition, we refer readers to [Saf22].

2.1 Basic Lattice Concepts

Recall from the introduction that for a linearly independent matrix $M \in \mathbb{R}^{n \times d}$, the lattice $\mathcal{L}[M]$ is the image of $\mathbb{Z}^d$ under M . It is often convenient to describe a lattice differently—as the set of integer solutions to a homogeneous linear system.

Definition 2.1 (Integer kernel). *For $A \in \mathbb{Z}^{m \times n}$, the integer kernel is*

$$\ker_{\mathbb{Z}}(A) \stackrel{\text{def}}{=} \{z \in \mathbb{Z}^n \mid Az = \mathbf{0}\},$$

which is a lattice in $\mathbb{Z}^n$ (and hence in $\mathbb{R}^n$).

One can efficiently convert $\ker_{\mathbb{Z}}(A)$ to a representation $\mathcal{L}[M]$ of the same lattice

Fact 2.2 (folklore). *Given $A \in \mathbb{Z}^{m \times n}$, one can compute in polynomial time a basis $M \in \mathbb{Z}^{n \times \dim(\ker A)}$ such that $\ker_{\mathbb{Z}}(A) = \mathcal{L}[M]$.*

2.2 Shortest Vector and Gap Problems

Let $\mathcal{L} \subseteq \mathbb{R}^n$ be a lattice. For any ℓ_p norm with $p \geq 1$, the *successive minima* are

$$\lambda_k^{(p)}(\mathcal{L}) \stackrel{\text{def}}{=} \inf \{ r > 0 \mid \dim(\text{span}\{v \in \mathcal{L} \mid \|v\|_p \leq r\}) \geq k \}.$$

In particular, $\lambda_1^{(p)}(\mathcal{L})$ is the length (in ℓ_p) of the shortest nonzero lattice vector.

Definition 2.3 (GapSVP). *Given a full-rank basis $M \in \mathbb{R}^{n \times d}$ and a threshold $C > 0$, the decision problem $\gamma\text{-GapSVP}_p$ asks to distinguish between:*

- **YES:** $\lambda_1^{(p)}(\mathcal{L}[M]) \leq C$.
- **NO:** $\lambda_1^{(p)}(\mathcal{L}[M]) > \gamma(n) \cdot C$.

Definition 2.4 (Unique-SVP; decision). *Given a full-rank basis $M \in \mathbb{R}^{n \times d}$ and a threshold $C > 0$, the decision problem $\gamma\text{-uSVP}_p$ asks to distinguish between:*

- **YES:** $\lambda_1^{(p)}(\mathcal{L}[M]) \leq C$ and $\lambda_2^{(p)}(\mathcal{L}[M]) \geq \gamma(n) \cdot C$.
- **NO:** $\lambda_1^{(p)}(\mathcal{L}[M]) > \gamma(n) \cdot C$.

Throughout the paper, we often abbreviate ℓ_p and $\gamma(n)$, writing simply **GapSVP** or **uSVP** when the parameters are clear from context. We also ignore floating-point precision, as it is insignificant.

2.3 Constraint Satisfaction Problems

Constraint Satisfaction Problems (CSPs) generalize problems with local consistency constraints. An important subcase is *Constraint Satisfaction Graph* (CSG), in which each constraint involves a pair of variables.

Definition 2.5 (*Constraint Satisfaction Graph*). *A CSG instance consists of a graph $G = (V, E)$, a finite alphabet Σ , and, for each edge $e \in E$, a constraint $\Phi_e \subseteq \Sigma \times \Sigma$.*

An assignment $c: V \rightarrow \Sigma$ satisfies the CSG instance if $(c(u), c(v)) \in \Phi_{(u,v)}$ for every edge $(u, v) \in E$. The decision problem is to determine whether a satisfiable assignment exists.

2.4 Promise-UP

The class UP (Unambiguous Non-deterministic Polynomial-Time) consists of decision problems solvable by a non-deterministic polynomial-time machine that has at most one accepting computation path for each input. Formally, a language L is in UP if there exists an efficient verifier V such that:

- If $x \in L$, then there exists a *unique* witness $w \in \{0, 1\}^*$ such that $V(x, w) = \mathbf{Yes}$. The witness w is of length polynomial in $|x|$.
- If $x \notin L$, then for all $w \in \{0, 1\}^*$, the verifier rejects, namely, $V(x, w) = \mathbf{No}$.

The class Promise-UP is the promise-problem analogue of UP. A problem is in Promise-UP if the YES instances have a single witness (and the NO instances have none). The difference is that some instances fall into neither the YES nor the NO cases.

An important example is **Unambiguous-3SAT**, the problem of deciding 3SAT instances with a promise of a unique satisfying assignment.

2.5 Finite fields

Let $\mathbb{E}$ be a field. A subset $\mathbb{F} \subseteq \mathbb{E}$ is a subfield if it is closed under the operations of $\mathbb{E}$ and forms a field with the induced operations. We write $\mathbb{E}/\mathbb{F}$ to denote that $\mathbb{E}$ is an extension of $\mathbb{F}$. Let $q = p^n$ be a prime power. We recall two well-known facts.

1. There exists, up to isomorphism, a unique field of size q , denoted $\mathbb{F}_q$.
2. For every integer $m > 1$, $\mathbb{F}_{q^m}/\mathbb{F}_q$.

It is well known that the extension $\mathbb{F}_{q^m}/\mathbb{F}_q$ can be constructed in time $\text{poly}(q^m)$.

2.6 Low-Degree Polynomials over Finite Fields

Let $\mathbb{F}$ be a finite field and consider variables $x_1, \dots, x_t$ over $\mathbb{F}$. A *monomial* is a product $x_1^{i_1} \cdots x_t^{i_t}$, with *total degree* $\deg(x_1^{i_1} \cdots x_t^{i_t}) \stackrel{\text{def}}{=} i_1 + \cdots + i_t$ and *individual degree* $\text{ideg}(x_1^{i_1} \cdots x_t^{i_t}) \stackrel{\text{def}}{=} \max\{i_1, \dots, i_t\}$. The total (resp. individual) degree of a polynomial is the maximum total (resp. individual) degree among its monomials. For any integer $d \geq 0$, define

$$\mathbb{F}_{\leq d}[x_1, \dots, x_t] = \{f : \mathbb{F}^t \rightarrow \mathbb{F} \mid f \text{ is a polynomial with } \deg(f) \leq d\},$$

and write $\mathbb{F}_{\leq d}^t$ when the variables are clear. We also work with *affine planes* in $\mathbb{F}^t$, namely two-dimensional affine subspaces, denoted by $\mathbf{PL}(\mathbb{F}^t)$. For a plane $\mathcal{P} \in \mathbf{PL}(\mathbb{F}^t)$, let

$$\mathcal{P}_{\leq d} = \{g : \mathcal{P} \rightarrow \mathbb{F} \mid \deg(g) \leq d\}.$$

Fact 2.6 (Low-Degree Extension). *Let $\mathbb{H} \subseteq \mathbb{F}$ and $f : \mathbb{H}^t \rightarrow \mathbb{F}$ be any function. There is a unique polynomial extension $f' : \mathbb{F}^t \rightarrow \mathbb{F}$ satisfying*

$$\forall x \in \mathbb{H}^t : f(x) = f'(x) \quad \text{and} \quad \text{ideg}(f') \leq |\mathbb{H}| - 1.$$

Proof. Follows immediately by interpolation. □

Lemma 2.7 (Schwartz–Zippel). *If $p \in \mathbb{F}_{\leq d}[x_1, \dots, x_t]$ is nonzero and $S \subseteq \mathbb{F}$, then*

$$\Pr_{r \in S^t} [p(r) = 0] \leq \frac{d}{|S|}.$$

2.7 Plane-vs-Plane

The Plane-vs-Plane test (introduced by Raz and Safra [RS97]) checks whether a purported encoding of a low-degree function is consistent.

Definition 2.8 (Plane-vs-Plane). *Let $\mathbb{F}$ be a finite field and let d be a positive integer. Given a table T that assigns to each affine plane $\mathcal{P} \in \mathbf{PL}(\mathbb{F}^t)$ a low-degree polynomial $T[\mathcal{P}] \in \mathcal{P}_{\leq d}$, the Plane-vs-Plane test proceeds as follows:*

1. *Pick a random affine line $\ell \subset \mathbb{F}^t$. Sample two distinct affine planes $\mathcal{P}_1$ and $\mathcal{P}_2$ containing ℓ (precisely, $\mathcal{P}_1 \cap \mathcal{P}_2 = \ell$).*
2. *Verify that the two functions agree on ℓ , namely, that $T[\mathcal{P}_1]|_{\ell} = T[\mathcal{P}_2]|_{\ell}$.*

Theorem 2.9 (Plane-vs-Plane test [RS97]). *Let $T : \mathbf{PL} \rightarrow \mathbb{F}_{\leq d}[x, y]$ be an assignment of one low-degree polynomial to each plane. There exists a constant $c > 0$, such that for every $\delta > 0$, if the test passes with probability δ , there exist a low-degree polynomial $g \in \mathbb{F}_{\leq d}[x_1, \dots, x_t]$ such that:*

$$\Pr_{\mathcal{P}} [T[\mathcal{P}] = g|_{\mathcal{P}}] \geq \delta - t \cdot \left(\frac{d}{|\mathbb{F}|} \right)^c$$

Here we require a different notion of soundness—*list-decoding soundness*. In the basic (unique-decoding) setting, a test passes with a non-negligible probability only if the local views partially agree with a single designated global function. In the list-decoding variant, we allow a short list of candidate global functions, and agreement with *any* one of them suffices. Although it was likely folklore that such statements follow from the same techniques, the earliest explicit formulation we know in this context is due to Moshkovitz and Raz [MR10], who also give a general recipe for deriving list-decoding guarantees from unique-decoding ones:

Theorem 2.10 (Plane-vs-Plane [RS97]: list-decoding). *Let $T: \mathbf{PL} \rightarrow \mathbb{F}_{\leq d}[x, y]$ be an assignment of one low-degree polynomial to each plane. There exists a constant $c > 0$, such that for every $\delta > 0$, there exists $k = O(\frac{1}{\delta})$ and a list of low-degree polynomials $g_1, \dots, g_k \in \mathbb{F}_{\leq d}[x_1, \dots, x_t]$ such that:*

$$\Pr_{P_1 \cap P_2 = \ell} \left[T[\mathcal{P}_1]_{|\ell} = T[\mathcal{P}_2]_{|\ell} \wedge \nexists i: (T[\mathcal{P}_1] = g_{i|P_1} \wedge T[\mathcal{P}_2] = g_{i|P_2}) \right] \leq \delta + t \cdot \left(\frac{d}{|\mathbb{F}|} \right)^c$$

In words: except with probability at most $\delta + t \cdot (d/|\mathbb{F}|)^c$, every passing test (two planes agreeing on their intersection line) is explained by one of the k global low-degree polynomials on the list.

2.7.1 Plane-vs-Plane, the Graph

It is often helpful to think of this test as sampling an edge in the appropriate graph whose vertices consist of all planes. We define the *Plane-vs-Plane graph* $G_{\mathcal{P}v\mathcal{P}} = (V, E)$ over $\mathbb{F}_q^t$:

1. Vertices: all affine 2-dimensional subspaces (planes) in $\mathbb{F}_q^t$.
2. Edges: an (undirected) edge connects two planes if their intersection is an affine line.

Observe that the $\mathcal{P}v\mathcal{P}$ graph is regular, that is, every vertex has the same degree.

For any subset $S \subseteq V$, we define the *edge expansion* of S as

$$\Phi(S) \stackrel{\text{def}}{=} \Pr_{(u,v) \in E} [v \notin S \mid u \in S]$$

that is, the probability a random edge coming out of S escapes to $V \setminus S$.

The proof of the 2-to-2 Games Theorem [KMS17, DKK⁺17, DKK⁺18, KMS23] was completed through an analysis of set expansion in the Grassmann graph. These expansion properties have since led to several applications, including improved low-degree testing [KM25] and new PCP theorems [MZ24]. For our purposes, weaker expansion results suffice (see Appendix A for a proof):

Fact 2.11. *Let $G_{\mathcal{P}v\mathcal{P}}$ be the Plane-vs-Plane graph over $\mathbb{F}_q^t$. Then for every subset $S \subseteq V$, we have*

$$\Phi(S) \geq 1 - \frac{|S|}{|V|} - \frac{3}{q}.$$

3 Auxiliary Techniques

In this section we present the techniques that are *special* to our reductions to SVP in the ℓ_p norm. Some ingredients are adapted from the PCP toolbox—in their algebraic version [DFK⁺11], introduced for lattices in [DKS98, DKRS03]. We tailor them to the ℓ_p regime. We present these components before formally describing the reductions, in order to clarify their role by separating the core ideas from the complications of the full reductions.

3.1 Super-Assignments

To construct SVP instances, we encode CSP via *super-assignments*, following the framework of [DKRS03].

Definition 3.1 (Super-Assignment). *Let $x_1, \dots, x_n$ be CSP variables taking values in the alphabet $\Sigma \stackrel{\text{def}}{=} \{\alpha_1, \dots, \alpha_k\}$. A super-assignment is an integer vector with an entry for each pair $(x_i, \alpha_j) \in \{x_1, \dots, x_n\} \times \Sigma$, namely,*

$$(x_i, \alpha_j) \mapsto v_{x_i, \alpha_j} \in \mathbb{Z}$$

Definition 3.2 (Natural Assignment). *Let $g: \{x_1, \dots, x_n\} \rightarrow \Sigma$ be an assignment to a CSP. The corresponding natural assignment is the super-assignment defined by:*

$$v_{x_i, \alpha_j} = \begin{cases} 1 & \text{if } \alpha_j = g(x_i), \\ 0 & \text{otherwise,} \end{cases}$$

this particular super-assignment is denoted $\langle g \rangle$.

Super-assignments are integer vectors, so one may impose homogeneous linear constraints on them to define a lattice. A core technique in [DKRS03] is using a low-degree test: Consider the Plane-vs-Plane test. For each plane $\mathcal{P} \in \mathbf{PL}(\mathbb{F}^t)$ and a low-degree polynomial $g \in \mathcal{P}_{\leq d}$ (potentially assigned to that plane), we introduce a variable $\mathcal{A}_{\mathcal{P}}[g] \in \mathbb{Z}$. Presuming a *global* degree- d polynomial $g \in \mathbb{F}_{\leq d}[x_1, \dots, x_t]$, the resulting natural assignment is as follows:

$$\langle g \rangle[\mathcal{P}, f] = \begin{cases} 1 & \text{if } f = g|_{\mathcal{P}}, \\ 0 & \text{otherwise.} \end{cases}$$

In our reductions, we will use similar variables. Using the Plane-vs-Plane test, we will describe homogeneous linear constraints that ensure all low-norm assignments are of the form $\alpha_1 \langle g_1 \rangle + \dots + \alpha_k \langle g_k \rangle$, for a small k . Additional constraints will encode the clauses of the original CSP instances.

3.2 The Advantage of ℓ_p Norms

After presenting super-assignments and the $\mathcal{PvP}$ graph, we now describe how they interact and why high-index ℓ_p norms are essential.

Definition 3.3. *Let $\mathcal{A}$ be a super-assignment for the $\mathcal{PvP}$ graph over $\mathbb{F}^t$. For each plane $\mathcal{P} \in \mathbf{PL}$, define its support as the set of functions that are assigned a nonzero value:*

$$\text{supp}[\mathcal{P}]_{\mathcal{A}} := \{f \in \mathcal{P}_{\leq d} \mid \mathcal{A}_{\mathcal{P}}[f] \neq 0\}.$$

When clear from context, we omit the subscript $\mathcal{A}$ and write $\text{supp}[\mathcal{P}]$.

For natural assignments, we have $|\text{supp}[\mathcal{P}]_{\mathcal{A}}| = 1$ for every $\mathcal{P} \in \mathbf{PL}$.

The size of the support is of interest because we can characterize $\mathcal{PvP}$ super-assignment whose support is bounded on every plane. More concretely, in Section 6.6 we prove that there exists a small constant $\varepsilon > 0$ such that the only super-assignments satisfying a certain set of linear constraints, for which $|\text{supp}[\mathcal{P}]| \leq |\mathbb{F}|^{\varepsilon}$ holds on all planes, are of the form $\mathcal{A} = a_1 \cdot \langle g_1 \rangle + \dots + a_k \cdot \langle g_k \rangle$, with $k \leq |\mathbb{F}|^{\varepsilon}$. This characterization plays a key role in the soundness analysis for Theorem 1.2 and Theorem 1.4

3.2.1 Rotation

To ensure that short vectors in the lattice, namely super-assignments with small norm, must have small supports, we add rotations. This technique was also used in [Din02] for the ℓ_{∞} norm.

Fact 3.4. *For every $x \in \mathbb{R}^m$ and $p \geq 2$, we have*

$$m^{\frac{1}{2} - \frac{1}{p}} \cdot \|x\|_p \geq \|x\|_2.$$

Proof. Hölder's inequality states $|\langle x, y \rangle| \leq \|x\|_p \|y\|_q$ for $\frac{1}{p} + \frac{1}{q} = 1$. Applying the inequality on $(x_1^2, \dots, x_m^2)$ and $(1, \dots, 1)$ with $\frac{2}{p} + \frac{p-2}{p} = 1$:

$$\|x\|_2^2 = \langle \vec{1}, (x_1^2, \dots, x_m^2) \rangle \leq m^{\frac{p-2}{p}} \cdot (\sum (x_i^2)^{p/2})^{2/p} = (m^{\frac{1}{2} - \frac{1}{p}} \|x\|_p)^2$$

Taking square roots gives $m^{\frac{1}{2} - \frac{1}{p}} \cdot \|x\|_p \geq \|x\|_2$. $\square$

Let $\mathcal{A}$ be a super-assignment and fix a plane $\mathcal{P} \in \mathbf{PL}$. Rotation preserves the ℓ_2 norm, so applying any rotation matrix U gives:

$$m^{\frac{1}{2} - \frac{1}{p}} \cdot \|U \cdot \mathcal{A}_{\mathcal{P}}\|_p \geq \|U \cdot \mathcal{A}_{\mathcal{P}}\|_2 = \|\mathcal{A}_{\mathcal{P}}\|_2 \geq \sqrt{|\text{supp}[\mathcal{P}]|}.$$

The key point is that we can construct a specific U that minimizes the ℓ_p norm of natural assignments.

Constructing the Rotation. Recall the recursive definition of the Hadamard matrix: $H_1 = (1)$, and for $n > 1$,

$$H_n = \begin{pmatrix} H_{n-1} & H_{n-1} \\ H_{n-1} & -H_{n-1} \end{pmatrix}.$$

We construct our rotation matrix from a normalized Hadamard matrix, removing some columns to accommodate general dimensions (not necessarily powers of 2).

Definition 3.5. For any $n \in \mathbb{N}$, let $m = 2^{\lceil \log n \rceil}$. Define the rotation matrix $U \in \mathbb{R}^{m \times n}$ to be a normalized submatrix of the Hadamard matrix:

$$U_{i,j} := \frac{1}{\sqrt{m}} H_{\log m}[i, j].$$

From now on, U denotes Definition 3.5 with proper dimensions. For each standard basis vector e_i , we have $\|U \cdot e_i\|_p = m^{\frac{1}{p} - \frac{1}{2}}$, attaining equality in Fact 3.4. Note that if $\mathcal{A}$ is a natural assignment, each $\mathcal{P} \in \mathbf{PL}$ is assigned a unit vector. The choice of U guarantees that for such $\mathcal{A}$:

$$m^{\frac{1}{2} - \frac{1}{p}} \cdot \|U \cdot \mathcal{A}_{\mathcal{P}}\|_p = \|U \cdot \mathcal{A}_{\mathcal{P}}\|_2 = \|\mathcal{A}_{\mathcal{P}}\|_2 = 1.$$

3.3 Composition-Recursion

We use the Composition-Recursion framework [AS98] to obtain quasi-polynomial reductions. Specifically, an algebraic version of Composition-Recursion, with modifications tailored to our setting. This algebraic Composition-Recursion originated in [DFK⁺11], to prove low-error PCP (Probabilistically Checkable Proofs) theorems. [DKRS03] utilized it to prove hardness of approximation for the Closest Vector Problem (CVP), showing inapproximability within a ratio of $n^{\frac{c_p}{\log \log n}}$, where c_p is a constant depending only on p . Informally, the Composition-Recursion of [DFK⁺11] consists of two alternating steps:

1. *Encode* a low-degree polynomial over $\mathbb{F}^t$ via its restrictions to subspaces of fixed dimension. Herein, the restrictions are all to affine planes (2-dimensional subspaces).
2. *Embed* these subspaces into a higher-dimensional vector space (typically $\mathbb{F}^t$), which substantially reduces the polynomial's degree.

This process continues until the degree reaches a threshold.

Contrary to CVP, for SVP we can enforce only homogeneous linear constraints. This degrades the soundness under Composition-Recursion. To overcome it, we introduce a new step to the Composition-Recursion:

3. *Extend* the domain of the low-degree polynomial, from $\mathbb{F}_q^t$ to $\mathbb{F}_{q^2}^t$. Since $\mathbb{F}_q^t \subseteq \mathbb{F}_{q^2}^t$, the polynomial can be naturally extended to $\mathbb{F}_{q^2}^t$ by interpreting its coefficients over the larger field $\mathbb{F}_{q^2}$.

Field extensions are the *key new ingredient* in our Composition-Recursion. By cautiously enlarging the base field along the Composition-Recursion, we prevent the soundness degradation that previous frameworks suffer from. This iterative process induces a forest structure as follows: the trees correspond to planes $\mathcal{P}_1 \in \mathbf{PL}(\mathbb{F}_q^t)$. The children of each plane correspond to planes $\mathcal{P}_2 \in \mathbf{PL}(\mathbb{F}_{q^2}^t)$, and so on. For convenience, we identify each vertex by the path from its root—a tuple $(\mathcal{P}_1, \dots, \mathcal{P}_r)$ where $\mathcal{P}_1 \subseteq \mathbb{F}_q^t, \mathcal{P}_2 \subseteq \mathbb{F}_{q^2}^t, \dots, \mathcal{P}_r \subseteq \mathbb{F}_{q^{2^{r-1}}}^t$.




Figure 1: Forest induced by Composition-Recursion. A node at depth r correspond to a path $(\mathcal{P}_1, \dots, \mathcal{P}_r)$ with $\mathcal{P}_j \in \mathbf{PL}(\mathbb{F}_{q^{2^{j-1}}}^t)$. Edges are the steps in our composition (embedding, extending the field and restricting to planes). Only a few representative children are drawn; actual branching is larger.

The rest of this section is dedicated to surveying the mathematical definitions and facts behind embeddings and field extensions.

3.4 Embedding

Assume arbitrary domains $\mathcal{X}, \mathcal{Y}$; an *embedding* is a structure-preserving map $E: \mathcal{X} \rightarrow \mathcal{Y}$. We describe herein an embedding of $\mathbb{F}^k$ into $\mathbb{F}^{i \cdot k}$. Let $c \in \mathbb{N}_+$,

$$E_c((\xi_1, \dots, \xi_k)) \stackrel{\text{def}}{=} (\xi_1, \xi_1^c, \dots, \xi_1^{c^{i-1}}, \dots, \xi_k, \dots, \xi_k^{c^{i-1}}).$$

This embedding drastically reduces the individual degree:

Fact 3.6. *Let $f: \mathbb{F}^k \rightarrow \mathbb{F}$ be a polynomial of individual degree $\text{ideg}(f) < c^i$. There exists a single polynomial $g: \mathbb{F}^{i \cdot k} \rightarrow \mathbb{F}$, of $\text{ideg}(g) < c$, so that*

$$\forall x \in \mathbb{F}^k: f(x) = g(E_c(x)).$$

Proof. Let $x_1^{t_1} \dots x_k^{t_k}$ be a monomial with $0 \leq t_1, \dots, t_k < c$. For each t_j there exist integers $0 \leq a_1^j, \dots, a_i^j < c$ such that $t_j = a_1^j + a_2^j \cdot c + \dots + a_i^j \cdot c^{i-1}$. We map the monomial to one from $\mathbb{F}[x_{1,1}, \dots, x_{k,i}]$,

$$\phi(x_1^{t_1} \dots x_k^{t_k}) \stackrel{\text{def}}{=} \prod_{m \in [i], n \in [k]} a_m^n x_{m,n}.$$

Extending ϕ linearly to polynomials yields an isomorphism between the polynomials of $\text{ideg} < c^i$ in k variables and those of $\text{ideg} < c$ in $i \cdot k$ variables. $f((\xi_1, \dots, \xi_k)) = \phi(f)(E_c((\xi_1, \dots, \xi_k)))$ so $\phi(f)$ is the unique polynomial of $\text{ideg} < c$ satisfying the requirement. $\square$

To embed an affine plane $\mathcal{P} \in \mathbf{PL}$, recall $\mathcal{P} = \vec{\alpha} + \text{span}(\vec{x}, \vec{y})$, which defines a natural isomorphism between the polynomials over $\mathbb{F}^2$ and over $\mathcal{P}$. Abusing notation, we write $E(x)$ when the plane and c are clear from context.

3.5 Field Extensions

Let $x_1^{i_1} \cdots x_t^{i_t} : \mathbb{F}_q^t \rightarrow \mathbb{F}_q$ be a monomial. It naturally extends to a monomial on $\mathbb{F}_{q^2}^t$. Similarly, a polynomial extends to $\mathbb{F}_{q^2}^t$ by extending each monomial. The soundness analysis requires “reversing” that extension.

Claim 3.7. *Let $f : \mathbb{F}_{q^2}^t \rightarrow \mathbb{F}_{q^2}$ be a low-degree polynomial with $\deg(f) \leq d$. If*

$$\Pr_{x_1, \dots, x_t \in \mathbb{F}_q} [f(x_1, \dots, x_t) \in \mathbb{F}_q] > \frac{d}{q}$$

then all the coefficients of f are members of $\mathbb{F}_q$. That is, $f|_{\mathbb{F}_q^t} : \mathbb{F}_q^t \rightarrow \mathbb{F}_q$ and has the same (individual and total) degree.

Proof. Fix $\alpha \in \mathbb{F}_{q^2} \setminus \mathbb{F}_q$. It is not hard to verify that:

$$\mathbb{F}_{q^2} = \{a_1 + a_2 \cdot \alpha \mid a_1, a_2 \in \mathbb{F}_q\}$$

Every monomial is of the form $c \cdot x_1^{i_1} \cdots x_t^{i_t}$, for some $c \in \mathbb{F}_{q^2}$, and there exist $a_1, a_2 \in \mathbb{F}_q$ such that $c = a_1 + a_2 \alpha$, so we write $f = g + \alpha h$, where the coefficients in every monomial of h, g are elements of $\mathbb{F}_q$.

Observe that for every $x \in \mathbb{F}_q^t$, the values $h(x), g(x)$ are inside $\mathbb{F}_q$. Thus, $f(x) \in \mathbb{F}_q \iff h(x) = 0$. If f satisfies $\Pr_{x \in \mathbb{F}_q^t} [f(x) \in \mathbb{F}_q] > \frac{d}{q}$, then $\Pr_{x \in \mathbb{F}_q^t} [h(x) = 0] > \frac{d}{q}$ and so the Schwartz-Zippel [Lemma 2.7](#) implies $h = 0$. All of f 's monomials have coefficients from $\mathbb{F}_q$ and the claim follows. $\square$

4 SUBEXP Hardness

This section presents a sub-exponential deterministic reduction to GapSVP and uSVP, achieving a constant inapproximability factor—proving [Theorem 1.1](#) and [Theorem 1.3](#).

4.1 Parameters and notation

Throughout this section, we fix an arbitrary norm index $p > 2$. Our starting point is a 3SAT formula $\Phi = \varphi_1 \wedge \cdots \wedge \varphi_m$ over variables $x_1, \dots, x_n$; the problem is deciding its satisfiability. The reduction outputs an instance of $(\sqrt{2} - o(1))\text{-GapSVP}_p$.

Basic parameters.

- Let $\mathbb{F}$ be a finite field with $|\mathbb{F}| \geq n^{\frac{3}{1-2/p}}$. (Any sufficiently large field of size $n^{\Theta(1)}$ works.)
- Choose an arbitrary subset $\mathbb{H} \subset \mathbb{F}$ of size $|\mathbb{H}| = \lceil n^{1/3} \rceil$. Since $|\mathbb{H}|^3 \geq n$, we define an injective mapping from each variable x_i to a unique point $y_i \in \mathbb{H}^3$; fix any such mapping $x_i \rightarrow y_i$.
- Set $d = |\mathbb{H}| - 1$. a 3SAT witness is a function $\{y_1, \dots, y_n\} \subseteq \mathbb{H}^3 \rightarrow \{0, 1\}$. By [Fact 2.6](#), any function $\mathbb{H}^3 \rightarrow \mathbb{F}$ admits a unique extension to $\mathbb{F}^3 \rightarrow \mathbb{F}$ of individual degree at most d .

Collections of planes. To achieve a norm index p close to 2 we will consider a sub-collection of all affine planes in $\mathbb{F}^3$. Namely, we take:

1. *Parallel planes.* Let $\text{Par}(\mathbb{F}^3) \subseteq \mathbf{PL}(\mathbb{F}^3)$ be the family of affine planes parallel to the coordinate axes.
2. *Clause planes.* For every clause involving variables x_i, x_j, x_k (regardless of their sign in the clause), choose an arbitrary affine plane $\mathcal{P}_{y_i, y_j, y_k} \in \mathbf{PL}(\mathbb{F}^3)$ that contains the three points y_i, y_j, y_k . Note that unless the points are collinear, this plane is unique. Denote

$$\mathbf{P}_{\text{sat}} \stackrel{\text{def}}{=} \{\mathcal{P}_{y_i, y_j, y_k} \mid \text{the clause mentions } x_i, x_j, x_k\}.$$

For convenience, set $\mathcal{R} \stackrel{\text{def}}{=} \mathbf{P}_{\text{sat}} \cup \text{Par}(\mathbb{F}^3)$, a union that will recur frequently below.

4.2 The Intermediate Lattice

As a first step in the reduction we construct an intermediate lattice $\mathcal{L}[M_I]$, by describing a system of homogeneous linear equations. By [Fact 2.2](#), the space of integer solutions of such a system spans a lattice $\mathcal{L}[M_I]$. The system includes a variable $\mathcal{A}_{\mathcal{P}}[f]$ for every plane $\mathcal{P} \in \mathcal{R}$ and every low-degree polynomial $f \in \mathcal{P}_{\leq 3d}$. Every table $\mathcal{A}$ must satisfy three types of linear constraints:

Simple. We require that the sum of the super-assignment on every plane $\mathcal{P} \in \mathcal{R}$ is the same. Equivalently, there exists a fixed global constant $\kappa \in \mathbb{Z}$, such that

$$\sum_{f \in \mathcal{P}_{\leq 3d}} \mathcal{A}_{\mathcal{P}}[f] = \kappa \quad \text{for all } \mathcal{P} \in \mathcal{R}.$$

To do so, fix any plane $\mathcal{P}_1 \in \mathcal{R}$, and for every other plane $\mathcal{P}_2 \in \mathcal{R}$ add the homogeneous equation

$$\sum_{f \in \mathcal{P}_1 \leq 3d} \mathcal{A}_{\mathcal{P}_1}[f] = \sum_{f \in \mathcal{P}_2 \leq 3d} \mathcal{A}_{\mathcal{P}_2}[f]. \quad (1)$$

Testing consistency. In addition, we add constraints that enforce consistency between adjacent planes. For every two planes $\mathcal{P}_1, \mathcal{P}_2 \in \mathcal{R}$ with a nonempty intersection, and a point $x \in \mathcal{P}_1 \cap \mathcal{P}_2$, add the homogeneous equations

$$\forall a \in \mathbb{F}: \sum_{f(x)=a} \mathcal{A}_{\mathcal{P}_1}[f] = \sum_{f(x)=a} \mathcal{A}_{\mathcal{P}_2}[f]. \quad (2)$$

3SAT Constraints. To enforce clause satisfaction, we introduce constraints on the variables associated with the planes in $\mathbf{P}_{\text{sat}}$. For each clause φ_i over variables $x_{\alpha_i}, x_{\beta_i}, x_{\gamma_i}$, consider the associated plane $\mathcal{P}_{y_{\alpha_i}, y_{\beta_i}, y_{\gamma_i}} \in \mathbf{P}_{\text{sat}}$. Let $f: \mathcal{P}_{y_{\alpha_i}, y_{\beta_i}, y_{\gamma_i}} \rightarrow \mathbb{F}$ be a polynomial of total degree at most $3d$. We eliminate any f that either does not represent a boolean assignment, or fails to satisfy the clause φ_i . This is done by imposing the constraint $\mathcal{A}_{\mathcal{P}}[f] = 0$ whenever:

1. At least one of the values $f(y_{\alpha_i}), f(y_{\beta_i}), f(y_{\gamma_i})$ is not in $\{0, 1\}$, or
2. The assignment $(x_{\alpha_i} = f(y_{\alpha_i}), x_{\beta_i} = f(y_{\beta_i}), x_{\gamma_i} = f(y_{\gamma_i}))$ does not satisfy the clause φ_i .

4.3 The GapSVP instance

Consider the intermediate lattice $\mathcal{L}[M_I]$, whose vectors correspond to tables $\mathcal{A}$ satisfying all constraints from the previous subsection. To amplify soundness, we multiply by a unitary matrix $\tilde{U}$ and define our final lattice as

$$M_F \stackrel{\text{def}}{=} (|\mathcal{R}|)^{-\frac{1}{p}} \cdot \tilde{U} \cdot M_I.$$

The output of this reduction is the GapSVP instance $(M_F, 1)$.

The matrix $\tilde{U}$ is constructed by placing copies of the rotation matrix U (from Definition 3.5) along the diagonal and normalizing. Namely, if $m \stackrel{\text{def}}{=} 2^{\lceil \log |\mathcal{P}_{\leq 3d}| \rceil}$ is the number of rows in each copy of U , the matrix $\tilde{U}$ is defined by

$$\tilde{U} \stackrel{\text{def}}{=} m^{\frac{1}{2} - \frac{1}{p}} \cdot \text{Diag}(U, \dots, U).$$

In words, we apply U to each $\mathcal{A}_{\mathcal{P}}$ and rescale by $m^{\frac{1}{2} - \frac{1}{p}}$.

Running time. The number of planes is $|\mathcal{R}| = n^{\Theta(1)}$, and each plane contributes $|\mathcal{P}_{\leq 3d}| = |\mathbb{F}|^{\binom{3d+2}{2}} = 2^{O(n^{2/3} \log n)}$ variables. Hence M_I is of size $2^{O(n^{2/3} \log n)}$. Multiplying by the $\tilde{U}$ increases the dimensions by at most a factor of 2, which does not change the asymptotics. Since the running time is polynomial in the lattice dimension, the reduction's time complexity is $2^{O(n^{2/3} \log n)}$.

Completeness. Let $\sigma: \{x_1, \dots, x_n\} \rightarrow \{0, 1\}$ be a satisfying assignment. Our global polynomial is the low-degree extension of $x_u \rightarrow \sigma(u)$, guaranteed to exist by Fact 2.6. We denote it $g: \mathbb{F}^3 \rightarrow \mathbb{F}$ (having $\text{ideg}(g) \leq d$).

Consider the natural assignment $\langle g \rangle$. Namely, for $\mathcal{P} \in \mathcal{R}$ and $f \in \mathcal{P}_{\leq 3d}$, we assign $\langle g \rangle[\mathcal{P}, f] = 1$ if $g|_{\mathcal{P}} = f$ and otherwise $\langle g \rangle[\mathcal{P}, f] = 0$. It can be seen that $\langle g \rangle$ satisfies the *simple*, *consistency*, and *3SAT* constraints. Thus, it is a vector in the intermediate lattice. For every $i \leq k$, it holds $\| \frac{m^{\frac{1}{2} - \frac{1}{p}}}{|\mathcal{R}|^{\frac{1}{p}}} \cdot U \cdot e_i \|_p = \frac{1}{|\mathcal{R}|^{\frac{1}{p}}}$. Hence, the vector $\frac{m^{\frac{1}{2} - \frac{1}{p}}}{|\mathcal{R}|^{\frac{1}{p}}} \text{Diag}(U, \dots, U) \cdot \langle g \rangle$ has an ℓ_p norm of exactly 1.

4.4 Soundness Analysis

Now, we prove that if Φ is an unsatisfiable 3SAT instance, then $\lambda_1^{(p)} \geq \sqrt{2} - \eta$ for every $\eta > 0$. Let $\eta > 0$ be an arbitrary constant, and assume, by way of contradiction, that

$$\lambda_1^p \left[\mathcal{L} \left[\frac{m^{\frac{1}{2} - \frac{1}{p}}}{|\mathcal{R}|^{\frac{1}{p}}} \text{Diag}(U, \dots, U) \cdot M \right] \right] < \sqrt{2} - \eta.$$

and fix a short nonzero vector $\frac{m^{\frac{1}{2} - \frac{1}{p}}}{|\mathcal{R}|^{\frac{1}{p}}} \text{Diag}(U, \dots, U) \cdot \mathcal{A}$. Denote $S \subseteq \mathcal{R}$, the subset of “bad planes”:

$$S \stackrel{\text{def}}{=} \{ \mathcal{P} \in \mathcal{R} \mid \forall i : \mathcal{A}_{\mathcal{P}} \neq \pm e_i \text{ and } \mathcal{A}_{\mathcal{P}} \neq 0 \}.$$

The soundness argument proceeds in several steps. We begin by establishing that S , the set of “bad planes”, is nonempty. Specifically, we show that if no plane is bad, then the underlying 3SAT formula is satisfiable. Next, we prove that S contains almost all the planes. We apply the Schwartz–Zippel Lemma 2.7 to show that nearly all neighbors of bad planes are themselves bad. Using the structure of the underlying graph, this property propagates, implying that the fraction of bad planes is $1 - o(1)$. Finally, we observe that planes in S have an increased norm. This leads to a contradiction, completing the argument.

Lemma 4.1. *S is nonempty.*

Proof. Recall that $\mathcal{A}$ satisfies the “simple” constraints (1), and therefore, there is a global constant $\kappa \in \mathbb{Z}$ such that all the planes satisfy $\sum_{g \in \mathcal{P}_{\leq 3d}} \mathcal{A}_{\mathcal{P}}[g] = \kappa$. Now, split into cases according to the value of κ :

- If $|\kappa| \neq 1$, S contains all planes with nonzero assignment. Because $\mathcal{A} \neq \vec{0}$, for at least one $\mathcal{P} \in \mathcal{R}$, $\mathcal{A}_{\mathcal{P}} \neq \vec{0}$ implying $\mathcal{P} \in S$.
- If $|\kappa| = 1$, we may consider $-\mathcal{A}$ instead, so w.l.o.g, $\kappa = 1$. Assume, by way of contradiction, that S is empty and let us construct a satisfying assignment to the 3SAT formula as follows:
 1. For a variable x_i , choose an arbitrary plane that contains its corresponding point $\mathcal{P} \ni y_i$.
 2. $\mathcal{A}_{\mathcal{P}}$ is a unit vector (since $S = \emptyset$); thus $\mathcal{A}_{\mathcal{P}}[g] = 1$ for a single g . We set $\sigma(x_i) \stackrel{\text{def}}{=} g(y_i)$.

The *consistency constraints* (2) imply that the assignment does not depend on the selection of $\mathcal{P}$. From the 3SAT’s constraints, the assignment satisfies the 3SAT formula. $\square$

We now present the main argument showing that S contains $1 - o(1)$ of the planes. The intuition is that for each $\mathcal{P} \in S$, either $\|\mathcal{A}_{\mathcal{P}}\|_p$ is large, or the anomalies propagate to neighboring planes via the *consistency constraints* (2).

Fix an arbitrary plane $\mathcal{P} \in S$. Since $\mathcal{A}$ is a short vector, we have $\|\frac{m^{\frac{1}{2}-\frac{1}{p}}}{|\mathcal{R}|^{\frac{1}{p}}} \cdot U \cdot \mathcal{A}_{\mathcal{P}}\|_p \leq \sqrt{2}$. [Fact 3.4](#) bounds the ratio between ℓ_2 and ℓ_p , resulting in $\|\frac{1}{|\mathcal{R}|^{\frac{1}{p}}} \cdot U \cdot \mathcal{A}_{\mathcal{P}}\|_2 \leq \sqrt{2}$. Rearranging gives $\|\mathcal{A}_{\mathcal{P}}\|_2 \leq \sqrt{2}|\mathcal{R}|^{\frac{1}{p}}$. Since $\mathcal{A}_{\mathcal{P}}$ is an integer vector, the number of nonzero coordinates, namely $|\text{supp}[\mathcal{P}]|$, is at most:

$$|\text{supp}[\mathcal{P}]| \leq \|\mathcal{A}_{\mathcal{P}}\|_2^2 \leq 2|\mathcal{R}|^{\frac{2}{p}}.$$

Fix $\mathcal{P} \in S$ and let $f \in \text{supp}[\mathcal{P}]$. For any distinct $g \in \text{supp}[\mathcal{P}]$, the functions f and g agree on at most $3d \cdot |\mathbb{F}|$ points (From the Schwartz-Zippel [Lemma 2.7](#), since they are different and $\deg f, \deg g \leq 3d$). Therefore, f disagrees with all of $\text{supp}[\mathcal{P}] \setminus \{f\}$ on a fraction of at least $\frac{|\mathbb{F}| - 3d \cdot |\text{supp}[\mathcal{P}]|}{|\mathbb{F}|}$ of the points in $\mathcal{P}$. Note that

$$|\mathcal{R}| \leq |\text{Par}(\mathbb{F}^3)| + |\mathbf{P}_{\text{sat}}| = 3|\mathbb{F}| + n^3 \leq 4|\mathbb{F}|.$$

Using the earlier bound on $|\text{supp}[\mathcal{P}]|$, f disagrees with all of $\text{supp}[\mathcal{P}] \cap \{f\}$ on a fraction of at least

$$1 - \frac{6d|\mathcal{R}|^{\frac{2}{p}}}{|\mathbb{F}|} \leq 1 - \frac{24(\lceil n^{1/3} \rceil - 1)|\mathbb{F}|^{\frac{2}{p}}}{|\mathbb{F}|} = 1 - 24(\lceil n^{1/3} \rceil - 1)(n^{\frac{3}{1-2/p}})^{p/2-1} = 1 - o(1).$$

Suppose $|\text{supp}[\mathcal{P}]| > 1$ and fix distinct $f, g \in \text{supp}[\mathcal{P}]$. For a $1 - o(1)$ fraction of the points $x \in \mathcal{P}$, f differs from every function in $\text{supp}[\mathcal{P}] \setminus \{f\}$, and g differs from every function in $\text{supp}[\mathcal{P}] \setminus \{g\}$. For every such $x \in \mathcal{P}$,

$$\sum_{h \in \mathcal{P}_{\leq 3d}, h(x)=f(x)} \mathcal{A}_{\mathcal{P}}[h] = \mathcal{A}_{\mathcal{P}}[f] \neq 0 \quad \text{and} \quad \sum_{h \in \mathcal{P}_{\leq 3d}, h(x)=g(x)} \mathcal{A}_{\mathcal{P}}[h] = \mathcal{A}_{\mathcal{P}}[g] \neq 0.$$

By the *consistency constraints* (2), any neighboring plane $\mathcal{P}' \in \mathcal{R}$ containing x has to satisfy the same equations. Hence, any $\mathcal{P}'$ containing x is also a “bad plane”.

If $|\text{supp}[\mathcal{P}]| = 1$, then $\mathcal{A}_{\mathcal{P}}$ has a single nonzero entry, with value different than ± 1 (since $\mathcal{P} \in S$). The *consistency constraints* (2) immediately imply that for every plane $\mathcal{P}' \in \mathcal{R}$ intersecting $\mathcal{P}$, it must also hold that $\mathcal{P}' \in S$.

We established that for every plane $\mathcal{P} \in S$, a $1 - o(1)$ fraction of the points $x \in \mathcal{P}$ are contained only in planes from S . Now, we show that S contains almost all of the planes.

First, observe that there exists $\mathcal{P} \in S \cap \text{Par}(\mathbb{F}^3)$. Recall:

$$\text{Par}(\mathbb{F}^3) = \bigcup_{a \in \mathbb{F}} \left\{ \{(a, y, z) \mid y, z \in \mathbb{F}\}, \{(x, a, z) \mid x, z \in \mathbb{F}\}, \{(x, y, a) \mid x, y \in \mathbb{F}\} \right\}.$$

The induced $\mathcal{PvP}$ graph on $\text{Par}(\mathbb{F}^3)$ contains three independent sets of sizes $|\mathbb{F}|$ (one for each axis), with all possible edges between different sets present. Since each “bad plane” has a $1 - o(1)$ fraction of its neighbors also in S , it follows that $(1 - o(1))|\text{Par}(\mathbb{F}^3)|$ planes are in S .

Therefore $|S| = (1 - o(1))|\text{Par}(\mathbb{F}^3)| = (1 - o(1))|\mathcal{R}|$. By applying [Fact 3.4](#) again, for every $\mathcal{P} \in S$,

$$\left\| \frac{m^{1/2-1/p}}{|\mathcal{R}|^{1/p}} \cdot U \cdot \mathcal{A}_{\mathcal{P}} \right\|_p \geq \frac{1}{|\mathcal{R}|^{1/p}} \cdot \|\mathcal{A}_{\mathcal{P}}\|_2 \geq \frac{\sqrt{2}}{|\mathcal{R}|^{1/p}}.$$

The last inequality holds because $\mathcal{P} \in S$, so it is not a unit vector or $\vec{0}$. Since S has a fractional size of $1 - o(1)$, we reach a contradiction.

4.5 Unique-SVP

The same reduction also establishes [Theorem 1.3](#). To this end, we start our reduction from Unambiguous-3SAT instances. Additionally, we assume the mapping $x_i \rightarrow y_i \in \mathbb{H}^3$ is bijective. If n is not a perfect cube, we may pad the formula with $o(n)$ dummy variables (constrained to be 0).

To enforce the uniqueness of the shortest vector, we introduce additional linear constraints. For every parallel plane $\mathcal{P} \in \text{Par}(\mathbb{F}^3)$ and function $f \in \mathcal{P}_{\leq 3d}$, we enforce

$$\text{If } \text{ideg}(f) > d, \quad \text{then } \mathcal{A}_{\mathcal{P}}[f] = 0.$$

We imposed restrictions on the reduction, so the soundness guarantee stays intact. The completeness is still easily verifiable, and it remains to prove uniqueness, i.e., $\lambda_2^{(p)} \geq \sqrt{2} - o(1)$. To do so, we show that all short vectors are of the form $\frac{m^{\frac{1}{2}-\frac{1}{p}}}{|\mathcal{R}|^{1/p}} \text{Diag}(U, \dots, U) \cdot \langle g \rangle$, and a single g exists for each satisfying assignment. For that purpose, fix any short vector.

The soundness analysis states that short vectors have no “bad” planes. Consequently, for every plane $\mathcal{P} \in \mathcal{R}$, $\mathcal{A}_{\mathcal{P}}$ is a unit vector—there is a unique function $f_{\mathcal{P}} \in \text{supp}[\mathcal{P}]$. Define a global function $G: \mathbb{F}^3 \rightarrow \mathbb{F}$ by selecting, for each $x \in \mathbb{F}^3$, a plane $\mathcal{P} \ni x$, and setting

$$G(x) \stackrel{\text{def}}{=} f_{\mathcal{P}}(x).$$

The *consistency constraints* (2) guarantee that G is well-defined (, independent of the choice of $\mathcal{P}$). By definition, $\mathcal{A} = \langle G \rangle$ is a natural assignment. The induced assignment $x_i \rightarrow G(y_i)$ satisfies the 3SAT formula, as the 3SAT constraints are enforced.

Since we started from an instance of Unambiguous-3SAT, the restriction $G|_{\mathbb{H}^3}$ is uniquely determined. [Fact 2.6](#) states that there is a unique low-degree extension of $x \in \mathbb{H}^3 \rightarrow G(x)$ on $\mathbb{F}^3$. It remains to show that $\text{ideg}(G) \leq d$. The additional constraints enforce that the restriction of G to each (affine) parallel plane $\mathcal{P} \in \text{Par}(\mathbb{F}^3)$ has individual degree at most d , and it is well-known that if all such restrictions have individual degree at most d , then $\text{ideg}(G) \leq d$, completing the proof.

5 The construction

We present a deterministic reduction from 3COL to GapSVP in ℓ_p . Formally, we prove:

Theorem 5.1 (Main; reduction). *Let $G = (V, E)$ be a 3COL instance with $n \stackrel{\text{def}}{=} |V|$. For every even integer $t \geq 4$, prime power $q = q(n) \geq n$, and $p \geq p_t > 2$, there is a deterministic reduction mapping G to an instance of GapSVP_{γ}^p on a lattice of dimension $n' = q^{O(\log^{1/\log(\frac{1}{2})} n)}$, with running time $q^{O(\log^{1/\log(\frac{1}{2})} n)}$ and gap ratio $\gamma = q^{1/p_t}$.*

[Theorem 1.2](#) follows by instantiating $q(\cdot)$ and t . If we set $q(n) \approx n$ and $t > 2^{1/\varepsilon+1}$, then $n' = n^{O(\log^\varepsilon n)}$ and, as a function of the output dimension n' , the approximation ratio becomes $2^{\Theta(\log n')^{\frac{1}{1+\varepsilon}}}$, and can be rewritten as $2^{\Theta(\log n')^{1-\varepsilon}}$. If instead $q(n) \approx 2^{n^\varepsilon}$, then $n' = 2^{O(n^\varepsilon \log^\varepsilon n)}$ and the ratio becomes $n'^{\Theta(1/\log \log n')^\varepsilon}$. The theorem holds for every $\varepsilon > 0$, allowing us to avoid asymptotic notations by tuning ε .

To prove hardness for unique instances ([Theorem 1.4](#)), minor adaptations are required, which we describe and analyze in [Section 6.5](#).

Basic parameters. Let $G = (V, E)$ be a 3COL instance. Most parameters match those of [Section 4](#).

- Let $t \geq 4$ be a fixed *even* integer. Throughout our construction, affine planes will be embedded into vector spaces of dimension t .
- Let $\mathbb{F} = \mathbb{F}_q$ denote a finite field of size $q \geq n = |V|$.
- Let $\mathbb{H} \subseteq \mathbb{F}$ be an arbitrary set with cardinality $|\mathbb{H}| = \lceil n^{\frac{1}{t}} \rceil$.
- Let $p_t > 2$ be the (minimal) norm parameter used in our soundness analysis. This is a fixed constant depending only on t ; The value of p_t is not explicitly stated and arises from assumptions throughout the soundness analysis. We prove soundness for ℓ_p norms where $p \geq p_t$.
- Again, we set $d \stackrel{\text{def}}{=} (|\mathbb{H}| - 1)$, the *individual* degree of low-degree extending a function on $\mathbb{H}^t$.

Identically to [Section 4](#), $|\mathbb{H}^t| \geq n$, so we injectively map each $v \in V$ to a unique $x_v \in \mathbb{H}^t$.

5.1 The Intermediate Lattice

Again, we define the intermediate lattice $\mathcal{L}[M_I]$ as the integer solutions of a system of homogeneous linear equations. Variables in this system correspond to the leaf nodes of the Composition-Recursion tree (affine planes over an extension field) and the low-degree polynomials over these planes.

5.1.1 Composition-Recursion Forest

We begin by specifying the degree bounds used in the Composition-Recursion forest. Except at the leaves, these bounds are not explicitly enforced by the construction; the soundness analysis will show that, for short vectors, the super-assignment associated with each subtree consists of only a few polynomials that satisfy the stated individual-degree bounds.

For root nodes (namely, affine planes $\mathcal{P} \in \mathbf{PL}(\mathbb{F}_q^t)$), we distinguish two cases. If $\mathcal{P}$ is axis-parallel, the bound is d . Otherwise, the bound is $t \cdot d$, since restricting a polynomial to an arbitrary affine plane preserves total degree but may increase the individual degree. We now define the bounds recursively. Suppose $\mathcal{P}_1 \subseteq \mathbb{F}_q^t$, $\mathcal{P}_2 \subseteq \mathbb{F}_{q^2}^t$, $\dots$, $\mathcal{P}_r \subseteq \mathbb{F}_{q^{2^{r-1}}}^t$ form a path from a root downward, and let d' be the bound for $(\mathcal{P}_1, \dots, \mathcal{P}_{r-1})$. Then the bound for $(\mathcal{P}_1, \dots, \mathcal{P}_r)$ is:

1. If $\mathcal{P}_r$ is axis-parallel, bound the degree by $\lfloor (d')^{2/t} \rfloor$.
2. Otherwise, bound the degree by $t \cdot \lfloor (d')^{2/t} \rfloor$.

Which corresponds to the decrease in ideg when embedding ([Fact 3.6](#)). This case distinction is necessary because [Fact 3.6](#) requires controlling the *individual* degree rather than the total degree. We iterate the Composition-Recursion and field extension until the individual-degree bound drops below $10t^2$.

At each leaf, if the path from the root is $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}$ and the current bound is d' , we introduce a variable for every function $f: \mathcal{P}_{r+1} \rightarrow \mathbb{F}_{q^{2^r}}$ with $\text{ideg}(f) \leq d'$. We denote this variable by $\mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid f]$.

5.1.2 Local-to-global constraints

Let $\mathcal{P}_1 \in \mathbf{PL}(\mathbb{F}_q^t), \dots, \mathcal{P}_k \in \mathbf{PL}(\mathbb{F}_{q^{2^k}}^t)$ be a (possibly empty) sequence of planes, and let d' denote the individual-degree bound at $\mathcal{P}_k$ (with $d' = d$ when $k = 0$). The subtree rooted at $(\mathcal{P}_1, \dots, \mathcal{P}_k)$ is intended to encode low-degree polynomials on $\mathcal{P}_k$ (or on $\mathbb{F}_q^t$ if $k = 0$).

We enforce a path-independence condition: for every $x \in \mathbb{F}_{q^{2^k}}^t$, if we embed x along any chain of planes down to a leaf, the assigned value at x is independent of the particular chain. Intuitively, the system encodes a low-degree polynomial by recursively encoding its restrictions to planes, and the constraints ensure that the evaluation at x is consistent across all plane choices.

To formalize this, fix any prefix $\mathcal{P}_1, \dots, \mathcal{P}_k$ and extend it along a path $\mathcal{P}_{k+1}, \dots, \mathcal{P}_{r+1}$ from $\mathcal{P}_k$ downward such that $x \in \mathcal{P}_{k+1}$, $E^{(1)}(x) \in \mathcal{P}_{k+2}$, $\dots$, $E^{(r-k)}(x) \in \mathcal{P}_{r+1}$ ¹. For any alternative extension $\tilde{\mathcal{P}}_{k+1}, \dots, \tilde{\mathcal{P}}_{r'+1}$ with $x \in \tilde{\mathcal{P}}_{k+1}$, $\dots$, $E^{(r'-k)}(x) \in \tilde{\mathcal{P}}_{r'+1}$, we impose the local-to-global consistency constraint

$$\forall a \in \mathbb{F}_{q^{2^k}} : \sum_{f(E^{(r-k)}(x))=a} \mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid f] = \sum_{f(E^{(r'-k)}(x))=a} \mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_k, \tilde{\mathcal{P}}_{k+1}, \dots, \tilde{\mathcal{P}}_{r'+1} \mid f]. \quad (3)$$

Moreover, to enforce descent to the parent subfield, if $a \notin \mathbb{F}_{q^{2^{k-1}}}$ then, for any f with $f(E^{(r-k)}(x)) = a$, we require

$$\mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_k, \tilde{\mathcal{P}}_{k+1}, \dots, \tilde{\mathcal{P}}_{r'+1} \mid f] = 0.$$

5.1.3 3COL constraints

Let $\{u, v\} \in E$ be an edge in the 3COL instance. We permit nonzero assignments only to low-degree functions with x_u and x_v properly colored. Suppose $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}$ is a path to a leaf such that

$$x_u, x_v \in \mathcal{P}_1, \dots, E^{(r)}(x_u), E^{(r)}(x_v) \in \mathcal{P}_{r+1}$$

We add the constraint $\mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid f] = 0$ whenever $f(E^{(r)}(x_u)) \notin \{0, 1, 2\}$ or $f(E^{(r)}(x_v)) \notin \{0, 1, 2\}$. To enforce proper coloring, we also add the constraints $\mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid f] = 0$ whenever $f(E^{(r)}(x_u)) = f(E^{(r)}(x_v))$.

5.2 The GapSVP instance

The final lattice constructed in [Section 4](#) is the result of multiplying M_I with the rotation matrix of [Definition 3.5](#). Herein, we apply the same trick on the lowest layer of the recursion. In contrast to [Section 4](#), planes are weighted according to their depth in the Composition-Recursion tree.

Rather than explicitly writing $M_F \stackrel{\text{def}}{=} \tilde{U} \cdot M_I$, we describe the operation of $\tilde{U}$ on vectors in the intermediate lattice. Denote $C_{r+1} \stackrel{\text{def}}{=} (|\mathbf{PL}(\mathbb{F}_q^t)| \cdots |\mathbf{PL}(\mathbb{F}_{q^{2^r}}^t)|)^{-\frac{1}{p}}$. Let $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}$ be a path to a leaf and U a rotation matrix from [Definition 3.5](#). For the final lattice, we add:

$$C_{r+1} \cdot m^{\frac{1}{2} - \frac{1}{p}} \cdot U \cdot \mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid *]$$

where m is the number of rows in U . The number of columns in U corresponds to the number of low-degree functions on $\mathcal{P}_{r+1}$, and the number of rows is the nearest larger power of 2.

5.2.1 Size of the Reduction

At each level of the Composition-Recursion tree, the individual degree d decreases to $d' \leq t \cdot \lfloor d^{2/t} \rfloor$. The process stops once the degree is below some constant. Let R be the depth of the Composition-Recursion tree. Writing $\log(d') \leq \log(t) + \frac{2}{t} \log(d)$, we obtain

$$R = \frac{\log \log d}{\log(t/2)} + O(1).$$

¹Here $E^{(j)}(x)$ denotes the image of x after j successive embeddings: $E^{(1)}(x)$ is the image after embedding into $\mathcal{P}_{k+1}$, $E^{(2)}(x)$ after embedding into $\mathcal{P}_{k+2}$, and so on.

Number of leaves. At depth i the field size is q^{2^i} , and the number of affine planes in $\mathbb{F}_{q^{2^i}}^t$ is at most $|\text{PL}(\mathbb{F}_{q^{2^i}}^t)| \leq (q^{2^i})^{3t}$. The total number of leaves is bounded by

$$\prod_{i=0}^{R+1} |\text{PL}(\mathbb{F}_{q^{2^i}}^t)| \leq \prod_{i=0}^{R+1} (q^{2^i})^{3t} = q^{3t(2^{R+2}-1)}.$$

Applying our bound on the depth of the recursion, $2^R = 2^{\frac{\log \log d}{\log(t/2)} + O(1)} = O((\log d)^{1/\log(\frac{t}{2})})$. Therefore, the number of leaves is at most $q^{O(\log^{1/\log(\frac{t}{2})}(d))}$.

Per-leaf. Each leaf plane stores an entry for every low-degree function on that plane. At the leaves, the individual degree is at most $10t^2$, so the number of entries is at most

$$|\mathbb{F}_{q^{2^{R+1}}}^{100t^4}| = q^{O(2^{R+1})} = q^{O(\log^{1/\log(\frac{t}{2})}(d))}.$$

Summing over the leaves, the number of entries— $\mathcal{L}[M_I]$ lattice’s dimension—is at most $q^{O(\log^{1/\log(\frac{t}{2})}(d))}$. Note that the dimension of the final lattice $\mathcal{L}[M_F]$ is at most doubled, which does not change the asymptotics. The time complexity is polynomial in the size of the output.

5.3 Completeness

Fix a satisfying coloring $\sigma: V \rightarrow \{0, 1, 2\}$, and let $g: \mathbb{F}^t \rightarrow \mathbb{F}$ be a low-degree extension of the map $x_v \rightarrow \sigma(v)$. By [Fact 2.6](#), the individual degree satisfies $\text{ideg}(g) \leq d$. Restricting, embedding, and extending g along the Composition-Recursion forest naturally defines a vector in $\mathcal{L}[M_I]$:

1. At the root of the Composition-Recursion forest, each vertex corresponds to an affine plane $\mathcal{P} \in \mathbf{PL}(\mathbb{F}_q^t)$. We restrict g to that plane, obtaining $g|_{\mathcal{P}}$.
2. If a vertex $(\mathcal{P}_1, \dots, \mathcal{P}_r)$ is associated with a polynomial f , we embed f according to [Fact 3.6](#), obtaining f^* . The low-degree polynomial f^* extends naturally when passing to an extension field. The child vertex $(\mathcal{P}_1, \dots, \mathcal{P}_r, \mathcal{P}_{r+1})$ is then assigned $f^*|_{\mathcal{P}_{r+1}}$.
3. At a leaf $(\mathcal{P}_1, \dots, \mathcal{P}_k)$ with a function f , we set $\mathcal{A}[(\mathcal{P}_1, \dots, \mathcal{P}_k) \mid h] = 1_f(h)$, exactly as in [Section 4](#).

Denote such an assignment $\langle\langle g \rangle\rangle$. It is not hard to confirm that $\langle\langle g \rangle\rangle$ satisfies the linear constraints, and thus $\langle\langle g \rangle\rangle \in \mathcal{L}[M_I]$. From now on, we also refer to $\langle\langle g \rangle\rangle$ as a natural-assignment. $\tilde{U} \cdot \langle\langle g \rangle\rangle \in \mathcal{L}[M_F]$ is a lattice vector of ℓ_p norm exactly 1.

6 Soundness

Fix a norm parameter $p \geq p_t$. Suppose $\tilde{U} \cdot \mathcal{A} \in \mathcal{L}[M_F]$ is a lattice vector with $\|\tilde{U} \cdot \mathcal{A}\|_p \leq q^{1/p_t}$. We will show that, for a sufficiently large absolute constant $p_t > 2$ (depending only on the construction parameters, e.g., t), the underlying 3COL instance must be satisfiable. For readability, we do not attempt to optimize p_t ; any fixed sufficiently large choice of p_t suffices.

Before diving into the soundness analysis, we first present the necessary analytic tools.

Definition 6.1 (Weak Plane-vs-Plane constraints). *These are “local-to-global” constraints implicitly enforced in [Section 5](#). Let $\mathcal{A} \neq \vec{0}$ be a super-assignment — a vector with an entry for each pair $\mathcal{P} \in \mathbf{PL}(\mathbb{F}^t)$ and $f \in \mathcal{P}_{\leq d}$. For every affine line $\ell = \mathcal{P}_1 \cap \mathcal{P}_2$, we add the equations:*

$$\forall x \in \ell, \forall a \in \mathbb{F}: \sum_{f(x)=a} \mathcal{A}_{\mathcal{P}_1}[f] = \sum_{f(x)=a} \mathcal{A}_{\mathcal{P}_2}[f]. \quad (4)$$

Thus, instead of enforcing consistency on the entire line ℓ , we enforce it pointwise at each $x \in \ell$.

Lemma 6.3 (*PvP local to global*). There exists an absolute constant $\varepsilon > 0$ such that the following holds. Let $\mathcal{A}$ be a super-assignment that assigns an integer to each $\mathcal{P} \in \mathbf{PL}$ and $f \in \mathcal{P}_{\leq d}$. Assume $d \leq |\mathbb{F}|^\varepsilon$ and $|\text{supp}[\mathcal{P}]| \leq |\mathbb{F}|^\varepsilon$ for every $\mathcal{P} \in \mathbf{PL}$.

If $\mathcal{A}$ satisfies the weak Plane-vs-Plane constraints, then there exist integers $a_1, \dots, a_k \in \mathbb{Z}$, $k \leq |\mathbb{F}|^\varepsilon$, and global degree- d functions $g_1, \dots, g_k: \mathbb{F}^t \rightarrow \mathbb{F}$ such that

$$\mathcal{A} = a_1 \cdot \langle g_1 \rangle + \dots + a_k \cdot \langle g_k \rangle.$$

Lemma 6.3 is proven in Section 6.6. By combining the field-extension step with high ℓ_p norms, we apply the lemma at the lowest recursion level. Working upwards along the Composition-Recursion tree, we conclude that any short vector is an integer combination of few natural assignments. Finally, we show that the resulting low-degree polynomials encode proper 3-colorings of the original 3COL instance.

6.1 Characterizing short vectors

Informally, our argument is as follows. At each level of the recursion, the *field-extension* step preserves a polynomial ratio between the number of vertices and the ambient field size. Consequently, any attempt to bypass the local-to-global constraints within a subtree must “spend mass”: it increases the ℓ_p norm on that subtree by a factor polynomial in the field size. Because we work with $p \geq p_t$ and p_t is a fixed large constant, these polynomial losses are amplified, so such *local* assignments necessarily have significantly larger ℓ_p norm.

In the next section, our goal is to prove an analogue of Lemma 6.3, on the whole Composition-Recursion tree — that $\mathcal{A}$ admits a representation by a small list $(a_i, f_i)_{i \in I}$. Formally, we aim to write $\mathcal{A} = \sum_{i \in I} a_i \langle f_i \rangle$, when $|I|$ is bounded. $\sum_{i \in I} a_i \langle f_i \rangle$ will also be referred to as a *super-assignment*.

Consistent vertices. We construct such a representation by working from the leaves upward. A vertex $\mathcal{P}_1, \dots, \mathcal{P}_r$, not necessarily a leaf, will be called *consistent* if there exists a list $(a_i, f_i)_{i \in I}$ of low-degree polynomials $\mathbb{F}_{q^{2^r}}^t \rightarrow \mathbb{F}_{q^{2^r}}$ (With the individual-degree satisfying the Composition-Recursion degree bound) such that:

1. $|I| \leq |\mathbb{F}_{q^{2^r}}|^{\frac{12-t}{p_t}}$, and
2. The assignment induced on the subtree agrees with $(a_i, f_i)_{i \in I}$.

Note that consistency propagates downward: if a vertex is consistent, then every descendant is also consistent. Our objective is therefore to prove that the root ($r = 0$) is consistent. For the sake of contradiction, assume otherwise, and let $\mathcal{P}_1, \dots, \mathcal{P}_r$ be an inconsistent vertex of maximal depth.

6.2 The offspring’s support

For convenience, assume that $\mathcal{P}_1, \dots, \mathcal{P}_r$ is not a leaf. The case of leaves is simpler and follows the same steps.

As $\mathcal{P}_1, \dots, \mathcal{P}_r$ is of maximal depth, each offspring $\mathcal{P}_1, \dots, \mathcal{P}_r, \mathcal{P}_{r+1}$ is explained by a short list $I = (a_i, f_i)$. Using the consistency of $\mathcal{P}_1, \dots, \mathcal{P}_r, \mathcal{P}_{r+1}$, and that $\tilde{U} \cdot \mathcal{A}$ has low norm, we improve the bound on (a_i, f_i) ’s length.

Let $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}, \dots, \mathcal{P}_k$ be a leaf. Again, $\text{supp}[\mathcal{P}_1, \dots, \mathcal{P}_k]$ contains the low-degree polynomials with nonzero coefficient. Similarly to Section 4, the support’s size gives a lower bound on the norm:

$$\begin{aligned} \|C_k \cdot U \cdot \mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_k \mid *]\|_2 &\leq \|C_k \cdot m^{\frac{1}{2} - \frac{1}{p}} \cdot U \cdot \mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_k \mid *]\|_p \\ C_k |\text{supp}[\mathcal{P}_1, \dots, \mathcal{P}_k]|^{\frac{1}{2}} &\leq \|C_k \cdot m^{\frac{1}{2} - \frac{1}{p}} \cdot U \cdot \mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_k \mid *]\|_p \end{aligned}$$

Returning to $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}$, the Schwartz-Zippel Lemma 2.7 implies that (f_i) collides on a fraction of at most $d |\mathbb{F}_{q^{2^{r+1}}}|^{\frac{24-t}{p_t} - 1}$ of the planes. p_t is sufficiently large, making the collisions negligible.

By definition, an offspring $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}, \mathcal{P}$ is specified by the list (a_i, f_i^*) , where f_i^* is obtained by embedding and extending $f_i|_{\mathcal{P}}$. Whenever (f_i) does not collide on $\mathcal{P}$, the list (a_i, f_i^*) contains $|I|$ distinct low-degree functions. This property persists downwards, yielding $|I|$ distinct functions for all but an $o(1)$ fraction of the descendants of $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}$.

Consequently, the norm on the descendants of $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}$ is at least:

$$(1 - o(1)) \cdot \left(\sum_{\mathcal{P}_{r+2}, \dots, \mathcal{P}_k} C_k^p |I|^{\frac{p}{2}} \right)^{\frac{1}{p}} = (1 - o(1)) \cdot C_{r+1} |I|^{\frac{1}{2}}$$

where the equality follows from $C_k^p = (|\mathbf{PL}(\mathbb{F}_q^t)| \cdots |\mathbf{PL}(\mathbb{F}_{q^{2^{k-1}}}^t)|)^{-1}$. The total norm is at most $q^{\frac{1}{p_t}}$, and thus $(1 - o(1)) \cdot C_{r+1} |I|^{\frac{1}{2}} \leq q^{\frac{1}{p_t}}$. A lower bound for C_{r+1} implies an upper bound on I . Using $|\mathbf{PL}(\mathbb{F}_x^t)| \leq x^{3t}$ and $p > p_t$, it holds that $C_{r+1} \geq q^{-\frac{3t}{p_t}} \cdots q^{-\frac{3t}{p_t} \cdot 2^r} = q^{-\frac{3t}{p_t} \cdot (2^{r+1} - 1)}$. we obtain:

$$(1 - o(1)) \cdot q^{-\frac{3t}{p_t} \cdot (2^{r+1} - 1)} |I|^{\frac{1}{2}} \leq q^{\frac{1}{p_t}}$$

$$|I|^{\frac{1}{2}} \leq q^{\frac{3t}{p_t} \cdot (2^{r+1})} = |\mathbb{F}_{q^{2^{r+1}}}^t|^{\frac{3t}{p_t}}$$

Exactly a square root of the requirement on $|I|$. Restating, $|I| \leq |\mathbb{F}_{q^{2^r}}^t|^{\frac{12 \cdot t}{p_t}}$.

6.3 Pulling up

Every offspring $\mathcal{P}_1, \dots, \mathcal{P}_r, \mathcal{P}$ has a short description $(a_i^{\mathcal{P}}, f_i^{\mathcal{P}})_{i \in I_{\mathcal{P}}}$, where $I_{\mathcal{P}} \leq |\mathbb{F}_{q^{2^r}}^t|^{\frac{12 \cdot t}{p_t}}$. To describe $\mathcal{P}_1, \dots, \mathcal{P}_r$, one may be tempted to use [Lemma 6.3](#). Indeed, this is our end goal. However, $(a_i^{\mathcal{P}}, f_i^{\mathcal{P}})$ aren't assignments to $\mathcal{P}$, but low-degree polynomials on $\mathbb{F}_{q^{2^{r+1}}}$ ($\mathcal{P}$ after embedding and extending the field). Before applying the lemma, it is necessary to “reverse” the field-extension and embedding.

Reversing the Composition-Recursion. We show that for every $f_i^{\mathcal{P}}$, the restriction $f_i^{\mathcal{P}}|_{\mathbb{F}_{q^{2^r}}^t}$ defines a function $\mathbb{F}_{q^{2^r}}^t \rightarrow \mathbb{F}_{q^{2^r}}$. By [Claim 3.7](#), it is sufficient to prove that $f_i^{\mathcal{P}}$ output values are in $\mathbb{F}_{q^{2^r}}$ with non-negligible probability when restricted to $\mathbb{F}_{q^{2^r}}^t$. If $(f_i^{\mathcal{P}})$ does not collide on a point $x \in \mathbb{F}_{q^{2^r}}^t$, then by the local-to-global constraints for all $i \in I_{\mathcal{P}}$ we have $f_i^{\mathcal{P}}(x) \in \mathbb{F}_{q^{2^{r+1}}}$. The Schwartz-Zippel [Lemma 2.7](#) ensures that such collisions rarely occur. Thus, [Claim 3.7](#) guarantees that the restriction $f_i^{\mathcal{P}}|_{\mathbb{F}_{q^{2^r}}^t}$ is itself a low-degree polynomial $\mathbb{F}_{q^{2^r}}^t \rightarrow \mathbb{F}_{q^{2^r}}$. Since $f_i^{\mathcal{P}}$ is a low-degree polynomial, it follows that $f_i^{\mathcal{P}}$ is exactly the low-degree extension of $f_i^{\mathcal{P}}|_{\mathbb{F}_{q^{2^r}}^t}$.

Using [Fact 3.6](#), for every $f_i^{\mathcal{P}}$ there exists a unique $g_i^{\mathcal{P}} : \mathcal{P} \rightarrow \mathbb{F}_{q^{2^r}}$ such that embedding $g_i^{\mathcal{P}}$ yields $f_i^{\mathcal{P}}|_{\mathbb{F}_{q^{2^r}}^t}$, and extending the field recovers $f_i^{\mathcal{P}}$.

Constructing a super-assignment. Now, the conditions for [Lemma 6.3](#) are satisfied. Consider the $\mathcal{P}\mathcal{V}\mathcal{P}$ graph on $\mathbb{F}_{q^{2^r}}^t$. For every plane $\mathcal{P}$, we assign $a_i^{\mathcal{P}}$ to $g_i^{\mathcal{P}}$, and 0 elsewhere.

By [Equation \(3\)](#), the weak Plane-vs-Plane constraints [\(4\)](#) are satisfied. Applying [Lemma 6.3](#), we conclude that there exists a global description $(a_i, f_i)_{i \in I}$, with $|I| \leq |\mathbb{F}_{q^{2^r}}^t|^{\frac{12 \cdot t}{p_t}}$.

Cautious readers may notice we are not yet finished. It remains to show that the functions (f_i) have low *individual* degree, matching the degree bounds prescribed by the Composition-Recursion forest. However, [Lemma 6.3](#) only guarantees that the (total) degree is low enough.

Recall that on *parallel* planes $\mathcal{P} \in \mathbf{PL}(\mathbb{F}_{q^{2^r}})$, the individual degree $\text{ideg}(g_i^{\mathcal{P}})$, is sufficiently low. By the Schwartz-Zippel [Lemma 2.7](#), the functions (f_i) collide on only a negligible fraction of these parallel planes. Thus, for every f_i , restriction to almost all parallel planes has low individual degree. It then follows that f_i has a low individual degree.

6.4 Satisfying the 3COL

Having established a global description (a_i, f_i) , we now argue that every f_i induces a satisfying assignment $u \rightarrow f_i(x_u)$ to the underlying 3COL instance.

Fix any point $x \in \mathbb{F}_q^t$ such that (f_i) does not collide on x , and let $\{u, v\} \in E$ be an edge. Let $\mathcal{P}_1$ be an affine plane containing x, x_u, x_v . Let $\mathcal{P}_2$ be an affine plane containing $E(x), E(x_u), E(x_v)$, the images of these points under the embedding of $\mathcal{P}_1$. We can continue recursively and obtain a path to a leaf $\mathcal{P}_1, \dots, \mathcal{P}_r$, together with corresponding points $\tilde{x}, \tilde{x}_u, \tilde{x}_v \in \mathcal{P}_r$.

Since for all $i \neq j$, it holds that $f_i(x) \neq f_j(x)$, recursively embedding and extending (f_i) yields a collection (f_i^*) of distinct low-degree functions on $\mathcal{P}_r$. For each f_i with nonzero coefficient, Φ 's constraints ensure that $(f_i^*(\tilde{x}_u), f_i^*(\tilde{x}_v))$ satisfies the 3COL constraint on $\{u, v\}$. $f_i(x_u) = f_i^*(\tilde{x}_u)$ and $f_i(x_v) = f_i^*(\tilde{x}_v)$, thus $u \rightarrow f_i(x_u)$ is a satisfying coloring.

6.5 Unique-SVP

In the soundness analysis, we showed that short vectors correspond to $\sum_{i \in I} a_i \langle f_i \rangle$, where $\text{ideg}(f_i) \leq d$. In order to start from Unambiguous-3SAT, we need to map a formula $\Phi = (\varphi_1 \wedge \dots \wedge \varphi_m)$ over variables $x_1, \dots, x_n$ to $\mathbb{H}^t$ and replace the 3COL's constraints.

Fix any mapping $\{\varphi_1, \dots, \varphi_m\} \cup \{x_1, \dots, x_n\} \rightarrow \mathbb{H}^t$. Again, we assume it is bijective (we may pad the formula with $o(n)$ dummy variables constrained to be 0). Instead of the 3COL constraints, we enforce:

- *Alphabet*: Let φ_i be a formula, $x_{\varphi_i} \in \mathbb{H}$ its corresponding point, and $\tilde{x}_{\varphi_i}$ the result of embedding x_{φ_i} until reaching a leaf $\mathcal{P}_1, \dots, \mathcal{P}_r$. We enforce $\mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid f] = 0$ whenever $f(\tilde{x}_{\varphi_i}) \notin \{1, \dots, 7\}$. Similarly, let x_i be a variable, $x_{x_i} \in \mathbb{H}^t$ its corresponding point, and $\tilde{x}_{x_i}$ the result of embedding. We enforce $\mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid f] = 0$ whenever $f(\tilde{x}_{x_i}) \notin \{0, 1\}$.
- *Consistency*: Let φ_i be a formula and x_{x_j} be a variable in φ_i . After embedding both until reaching a leaf $\mathcal{P}_1, \dots, \mathcal{P}_{r+1}$, we have $\tilde{x}_{\varphi_i}, \tilde{x}_{x_j} \in \mathcal{P}_{r+1}$. Thinking about $f(\tilde{x}_{\varphi_i})$ as 3 bits—each equal to 1 if the corresponding literal is satisfied, and $f(\tilde{x}_{x_j})$ as the assignment to x_j ; for every f on $\mathcal{P}_{r+1}$ we enforce $\mathcal{A}[\mathcal{P}_1, \dots, \mathcal{P}_{r+1} \mid f] = 0$ if these bits are inconsistent.

Proving that each $x_i \rightarrow f_i(x_i)$ is a satisfying assignment to Φ follows identically to 3COL (Section 6.4). Note that starting from 3SAT instead of 3COL would have been possible, but since the constraints are somewhat harder to follow, we chose to start from 3COL.

Now, we need to prove *uniqueness*, i.e., $\lambda_2^{(p)} \geq q^{1/p_i} \lambda_1^{(p)}$. Identically to Section 4.5, Fact 2.6 promises that each f_i is the unique low-degree extension of the single satisfying assignment, and so up to multiplication by a scalar, there exists a single short vector.

6.6 Characterizing short PVP super-assignments

To finalize the soundness analysis, it is left to prove Lemma 6.3. While [DKRS03] proved a similar result, we present a proof based on $G_{\mathcal{P}_v \mathcal{P}}$'s expansion, that leads to an arguably cleaner reduction. We prioritize readability over tight parameters; thus, some constants are not presented in their optimal form.

Expansion of nonzero planes: Before proceeding to Lemma 6.3, we claim assignments with bounded support have nonzero values on almost all the planes. Claim 6.2 allows us to ignore planes $\mathcal{P} \in \mathbf{PL}$, whenever $\mathcal{A}_{\mathcal{P}} = \vec{0}$.

Claim 6.2. *Let $\varepsilon > 0$ and $\mathcal{A} \neq \vec{0}$ be a vector over $\mathbb{Z}$, with an entry for every $\mathcal{P} \in \mathbf{PL}(\mathbb{F}^t)$ and $f \in \mathcal{P}_{\leq d}$. Suppose $\mathcal{A}$ satisfies weak Plane-vs-Plane (4) constraints, and for every plane $\mathcal{P} \in \mathbf{P}$, the support of $\mathcal{P}$ is bounded by $|\text{supp}[\mathcal{P}]| \leq |\mathbb{F}|^\varepsilon$. Then, $\mathcal{A}_{\mathcal{P}} = \vec{0}$ for a fraction of at most $(d+3)|\mathbb{F}|^{-1+\varepsilon}$ of the planes.*

Proof. Fix such $\mathcal{A}$ and let $S \subseteq \mathbf{PL}$ be the planes with nonzero assignments. Namely, $S \stackrel{\text{def}}{=} \left\{ \mathcal{P} \in \mathbf{PL} \mid \mathcal{A}_{\mathcal{P}} \neq \vec{0} \right\}$. We will prove it is poorly expanding, and apply [Fact 2.11](#) to show S contains almost all the planes.

Fix $\mathcal{P} \in S$ and an arbitrary function $f \in \text{supp}[\mathcal{P}]$. For any different function $g \in \text{supp}[\mathcal{P}]$, the Schwartz-Zippel [Lemma 2.7](#) states that f and g agree on a random point with probability at most $\frac{d}{|\mathbb{F}|}$. As $|\text{supp}[\mathcal{P}]| \leq |\mathbb{F}|^\varepsilon$, f disagrees with all of $\text{supp}[\mathcal{P}] \setminus \{f\}$ on all but a fraction of at most $\frac{d}{|\mathbb{F}|} \cdot |\mathbb{F}|^\varepsilon$ of the points.

Let $\mathcal{P}'$ be a neighbor of $\mathcal{P}$: $\mathcal{P}' \cap \mathcal{P} = \ell$. If there exists $x \in \ell$ such that $\forall g \in \text{supp}[\mathcal{P}] \setminus \{f\}: g(x) \neq f(x)$, then [\(4\)](#) implies $\mathcal{P}' \in S$. This happens with a probability of at least $1 - d \cdot |\mathbb{F}|^{\varepsilon-1}$, so $\Phi(S) \leq d \cdot |\mathbb{F}|^{\varepsilon-1}$. Applying [Fact 2.11](#) (on the expansion of the $\mathcal{P}v\mathcal{P}$ graphs):

$$d \cdot |\mathbb{F}|^{\varepsilon-1} \geq 1 - \frac{3}{|\mathbb{F}|} - \frac{|S|}{|\mathbf{PL}|}$$

$$(d+3) \cdot |\mathbb{F}|^{\varepsilon-1} \geq 1 - \frac{|S|}{|\mathbf{PL}|}$$

□

Lemma 6.3 (*PvP local to global*). *There exists an absolute constant $\varepsilon > 0$ such that the following holds. Let $\mathcal{A}$ be a super-assignment that assigns an integer to each $\mathcal{P} \in \mathbf{PL}$ and $f \in \mathcal{P}_{\leq d}$. Assume $d \leq |\mathbb{F}|^\varepsilon$ and $|\text{supp}[\mathcal{P}]| \leq |\mathbb{F}|^\varepsilon$ for every $\mathcal{P} \in \mathbf{PL}$.*

If $\mathcal{A}$ satisfies the weak Plane-vs-Plane constraints, then there exist integers $a_1, \dots, a_k \in \mathbb{Z}$, $k \leq |\mathbb{F}|^\varepsilon$, and global degree- d functions $g_1, \dots, g_k: \mathbb{F}^t \rightarrow \mathbb{F}$ such that

$$\mathcal{A} = a_1 \cdot \langle g_1 \rangle + \dots + a_k \cdot \langle g_k \rangle.$$

Proof. The proof consists of three stages:

1. First, find a small list of degree- d polynomials $g_1, \dots, g_k: \mathbb{F}^t \rightarrow \mathbb{F}$, so that almost all planes may be explained via restrictions of $\{g_i\}$. To do so, use the consistency of the Plane-vs-Plane test [\[RS97\]](#).
2. Then, use [Fact 2.11](#), to show that a large part of $\mathbf{PL}$ correspond super-assignment $\mathcal{A}' = a_1 \cdot \langle g_1 \rangle + \dots + a_k \cdot \langle g_k \rangle$.
3. Finally, apply [Claim 6.2](#) on $\mathcal{A} - \mathcal{A}'$, showing $\mathcal{A} - \mathcal{A}' = \vec{0}$.

Identifying consistent planes: A pair $(\mathcal{P}, \ell \in \mathcal{P})$ is “good”, if the assignment to $\mathcal{P}$ is nontrivial, and does not collide on ℓ . Formally, $\text{supp}[\mathcal{P}] \neq \emptyset$ and $\forall f \neq g \in \text{supp}[\mathcal{P}]: f|_\ell \neq g|_\ell$. An edge $\{\mathcal{P}_1, \mathcal{P}_2\}$ is good, if $(\mathcal{P}_1, \mathcal{P}_1 \cap \mathcal{P}_2)$ and $(\mathcal{P}_2, \mathcal{P}_1 \cap \mathcal{P}_2)$ are good. Finally, a plane $\mathcal{P} \in \mathbf{PL}$ is good if at least half the incident edges are good.

Observe that [Claim 6.2](#) states almost all the planes have nontrivial support. Moreover, the Schwartz-Zippel [Lemma 2.7](#), combined with $|\text{supp}[\mathcal{P}]| \leq |\mathbb{F}|^\varepsilon$, implies that for almost all $\ell \subseteq \mathcal{P}$, the pair $(\mathcal{P}, \ell)$ is good. From the union bound, $\{\mathcal{P}_1, \mathcal{P}_2\} \in E_{\mathcal{P}v\mathcal{P}}$ is almost always good — and therefore nearly all the planes are good.

Probabilistic setting: For every $\mathcal{P} \in \mathbf{PL}$, we sample $T[\mathcal{P}]: \mathcal{P}_{\leq d} \rightarrow \mathbb{F}$. If $\text{supp}[\mathcal{P}] \neq \emptyset$, we uniformly sample from $\text{supp}[\mathcal{P}]$. Otherwise, we sample a constant $T[\mathcal{P}] \in \mathcal{P}_{\leq 0}$.

Let $\mathcal{P} \in \mathbf{PL}$ be a good plane and $\mathcal{P} \cap \mathcal{P}' = \ell$ be a good edge. The support does not collide on ℓ so the linear constraints [\(4\)](#) imply a non-negligible agreement between the functions:

$$\forall f \in \text{supp}[\mathcal{P}], \forall x \in \ell: \exists g \in \text{supp}[\mathcal{P}']: f(x) = g(x)$$

The support of each plane is bounded by $|\mathbb{F}|^\varepsilon$, so each $f \in \text{supp}[\mathcal{P}]$ agrees with at least one $g \in \text{supp}[\mathcal{P}']$, on at least $\frac{1}{|\mathbb{F}|^\varepsilon}$ of the points in ℓ . From the Schwartz-Zippel [Lemma 2.7](#), for every $f \in \text{supp}[\mathcal{P}]$ there exists at least one $g \in \text{supp}[\mathcal{P}']$, such that $f|_\ell(x) = g|_\ell(x)$. The support's size is bounded, so for every $f \in \mathcal{P}_{\leq d}$ we have

$$\Pr_T [f|_{\mathcal{P} \cap \mathcal{P}'} = T[\mathcal{P}']|_{\mathcal{P} \cap \mathcal{P}'}] \geq |\mathbb{F}|^{-\varepsilon}.$$

We are interested in avoiding tables T with an abnormally low success probability of the PvP test over edges incident to $\mathcal{P}$. Luckily, this only happens in a negligible fraction of the tables:

$$\Pr_T \left[\Pr_{\mathcal{P}' \cap \mathcal{P} = \ell} [T[\mathcal{P}]|_\ell = T[\mathcal{P}']|_\ell] < \frac{1}{4} |\mathbb{F}|^{-\varepsilon} \right] = \mathbb{E}_{f \in \mathcal{P}_{\leq d}} \left[\Pr_T \left[\frac{1}{d_{\mathcal{P} \cap \mathcal{P}}} \sum_{\mathcal{P}' \cap \mathcal{P} = \ell} 1_{g|_\ell = T[\mathcal{P}']|_\ell} < \frac{1}{4} |\mathbb{F}|^{-\varepsilon} \mid T[\mathcal{P}] = f \right] \right] \leq$$

$\mathcal{P}$ is a good plane so at least $\frac{1}{2}$ of the edges are good. In addition, the indicators in the summation are independent and equal to 1 with a probability of at least $|\mathbb{F}|^{-\varepsilon}$, so we apply the well-known *Chernoff bound*:

$$\mathbb{E}_{f \in \mathcal{P}_{\leq d}} \left[\Pr_T \left[\frac{1}{2 \#(\text{good edges})} \sum_{\text{good } \mathcal{P}' \cap \mathcal{P} = \ell} 1_{f|_\ell = T[\mathcal{P}']|_\ell} < \frac{1}{4} |\mathbb{F}|^{-\varepsilon} \mid T[\mathcal{P}] = f \right] \right] \leq \left(\sqrt{\frac{2}{e}} \right)^{\#(\text{good edges})/|\mathbb{F}|^\varepsilon}$$

The number of good edges adjacent to $\mathcal{P}$ is much larger than $|\mathbb{F}|^\varepsilon$, so the probability is exponentially small. From now on, we ignore it and assume that for every good plane $\mathcal{P} \in \mathbf{PL}$:

$$\Pr_{\mathcal{P}' \cap \mathcal{P} = \ell} [T[\mathcal{P}]|_\ell = T[\mathcal{P}']|_\ell] \geq \frac{1}{4} |\mathbb{F}|^{-\varepsilon}.$$

Global list decoding: We use [\[RS97\] \(2.10\)](#) on T . ε is sufficiently small, so $\frac{1}{2} |\mathbb{F}|^{-3\varepsilon}$ is larger than the error term in the theorem. Thus, there exists $k = O(|\mathbb{F}|^{3\varepsilon})$ and $f_1, \dots, f_k: \mathbb{F}^t \rightarrow \mathbb{F}$, of degree d , such that:

$$\Pr_{\mathcal{P}_1 \cap \mathcal{P}_2 = \ell} [T[\mathcal{P}_1]|_\ell = T[\mathcal{P}_2]|_\ell \wedge \exists i: (T[\mathcal{P}_1] = f_i|_{\mathcal{P}_1} \wedge T[\mathcal{P}_2] = f_i|_{\mathcal{P}_2})] \leq |\mathbb{F}|^{-3\varepsilon}$$

For every good plane, the success probability is at least $\frac{1}{4} |\mathbb{F}|^{-\varepsilon}$. Thus, the entries of at least $1 - 8|\mathbb{F}|^{-2\varepsilon}$ of the good planes agree with some f_i (at least half the planes are good).

We sample $|\mathbb{F}|^{1.5\varepsilon}$ tables and consider their list-decodings. Applying the union bound, almost all the good planes agree with a function in every list-decoding. In addition, since $|\text{supp}[\mathcal{P}]| \leq |\mathbb{F}|^\varepsilon$, for almost all the planes, all the functions in $\text{supp}[\mathcal{P}]$ were used in at least one table. Denote by $f_1, \dots, f_{k'}$ the concatenation of all the list-decodings, $k' = O(|\mathbb{F}|^{4.5\varepsilon})$. Combining the previous claims, for almost all the planes $\mathcal{P} \in \mathbf{PL}$:

$$\forall g \in \text{supp}[\mathcal{P}] : \exists 1 \leq i \leq k' : f_i|_{\mathcal{P}} = g$$

And we denote the set of such planes, by $S_1 \subseteq \mathbf{PL}$.

Relating restrictions: The Schwartz-Zippel [Lemma 2.7](#) implies that two functions in $\{f_1, \dots, f_{k'}\}$ agree on a point with probability at most $\frac{d}{|\mathbb{F}|} \binom{k'}{2} = O(|\mathbb{F}|^{10\varepsilon-1})$. We assume ε is sufficiently small, so $\frac{d}{|\mathbb{F}|} \binom{k'}{2} = o(1)$. Thus, for almost all planes (and points), the restrictions of $\{f_i\}_{i=1}^{k'}$ are pairwise distinct. Let $S_2 \subseteq S_1$ be the subset of those planes in S_1 . Observe that S_2 also contains most of the planes.

The assignment for every plane in S_2 could be uniquely described with a super-assignment over $f_1, \dots, f_{k'}$ (the coefficient of f_i is $\mathcal{A}_{\mathcal{P}}[f_i|_{\mathcal{P}}]$). Denote $\pi: S_2 \rightarrow \mathbb{Z}^{k'}$ the mapping between a plane in S_2 , and the coefficients of that super-assignment. Consider the equivalence relation over S_2 : $\mathcal{P}_1 \sim \mathcal{P}_2 \iff \pi(\mathcal{P}_1) = \pi(\mathcal{P}_2)$. Next, we will show that few edges cross between different equivalence classes. Then, we deduce one class has to contain almost all the planes.

For a plane $\mathcal{P}_1 \in S_2$ and an edge $\{\mathcal{P}_1, \mathcal{P}_2\} \in E_{\mathcal{P}_v \mathcal{P}}$, we can classify the edge into three types:

1. *Bad* edges: $\mathcal{P}_2 \notin S_2$.
2. *Non-escaping* edges: $\mathcal{P}_2 \in [\mathcal{P}]_{/\sim}$.
3. *Crossing* edges: $\mathcal{P}_2 \not\sim \mathcal{P}_1$.

Since S_2 contains $1 - o(1)$ of the edges, and the graph is $d_{\mathcal{P}_v\mathcal{P}}$ -regular, the fraction of bad edges is $o(1)$. If $\{\mathcal{P}_1, \mathcal{P}_2\}$ is crossing, $f_1, \dots, f_{k'}$ has to collide on $\mathcal{P}_1 \cap \mathcal{P}_2$. Thus, the Schwartz-Zippel [Lemma 2.7](#) implies $o(1)$ of the edges are crossing—almost all edges are non-escaping. Due to [Fact 2.11](#), at least one equivalence class has a fractional size of $1 - o(1)$.

Endgame: Denote $\mathcal{A}' = a_1 \cdot \langle f_1 \rangle + \dots + a_{k'} \cdot \langle f_{k'} \rangle$, the super-assignment of that equivalence class. The vector $\mathcal{A} - \mathcal{A}'$ assigns $\vec{0}$ to the planes in that equivalence class—almost all the planes. From linearity, it satisfies the $\mathcal{P}_v\mathcal{P}$ constraints. The triangle inequality bounds the size of each plane’s support by $|\mathbb{F}|^\varepsilon + k'$. We choose a sufficiently small $\varepsilon > 0$, so [Claim 6.2](#) implies $\mathcal{A} - \mathcal{A}' = \vec{0}$.

We remark that k' is effectively bounded by $|\mathbb{F}|^\varepsilon$, because for planes where $f_1, \dots, f_{k'}$ don’t collide, the support’s size is $|\{f_i \mid a_i \neq 0\}|$ (if $a_i = 0$, f_i is meaningless and could be ignored). $\square$

7 Discussion and Open Problems

We extend the *super-assignment* framework of [[DFK⁺99](#), [Din02](#), [DKRS03](#)] to establish improved hardness-of-approximation for both GapSVP and uSVP. The main obstacle in porting from CVP to SVP is excluding *self-involved* super-assignments—those supported on only a small fraction of planes—whose artificially low norm does not reflect any global assignment. This stems from the homogeneity of SVP: a super-assignment that is zero on almost all planes cannot be ruled out naively by local constraints, unlike CVP where nontriviality can be enforced everywhere. Our first goal, therefore, is to ensure low-norm assignments *disperse* across the $\mathcal{P}_v\mathcal{P}$ graph.

Leveraging expansion in Grassmann graphs—a generalization of the plane-vs-plane graph—entered the PCP toolkit following its pivotal role in resolving the 2-to-2 Games Conjecture [[KMS17](#), [DKK⁺17](#), [DKK⁺18](#), [KMS23](#)]; see also Minzer’s thesis [[Min22](#)]. These results yield *structure theorems* for small sets of planes/-subspaces: unless such a set is extremely expanding, it must exhibit a rigid structure. This machinery has since been extended, leading to multiple applications [[KM25](#), [MZ23](#), [MZ24](#)]. In our setting, much weaker expansion properties suffice, though our approach was inspired by the new understanding of expansion in such graphs.

To reach quasi-polynomial instance size we adopt Composition–Recursion [[AS98](#)] in the algebraic variant of [[DFK⁺99](#), [DKRS03](#)]. Composition–Recursion creates super-polynomially many planes, causing self-involved assignments to reappear. We counter this by passing, at each recursion level, to an *extension field* of the current field, thereby keeping the number of planes/vertices polynomial in the field size at every level. Field extensions may be useful in further Composition–Recursion applications.

Starting from *unambiguous* problems is a natural starting point for our pipeline: SVP_p and $\alpha\text{-BDD}_p$ with $\alpha < \frac{1}{2}$ are themselves unique. Valiant–Vazirani show a *randomized* reduction from SAT to Unambiguous-3SAT and, more generally, a pathway from NP to Promise-UP [[VV85](#)]. We remark that deterministically reducing an NP-hard problem to uSVP would imply $\text{NP} = \text{Promise-UP}$.

7.1 Open problems

Several directions remain open. The central challenge is to obtain comparable (or even weaker) results in the Euclidean norm ℓ_2 .

Conjecture 7.1 (Toward ℓ_2). Under deterministic reductions, GapSVP is hard to approximate in the ℓ_2 norm—even the *exact* SVP problem is not known to be NP-hard (deterministically).

We conjecture that PCP-style techniques can be adapted to ℓ_2 . Self-involved assignments of low-norm in ℓ_2 appear highly structured, which makes their characterization and exclusion plausible. Even a sub-exponential derandomization for ℓ_2 would be compelling.

As emphasized in the introduction, uSVP underpins many cryptographic constructions; sharpening its complexity in ℓ_2 is therefore interesting. Our matching uSVP bounds strengthen the case for the difficulty of unique instances, yet the known ℓ_2 hardness still lags.

Conjecture 7.2 (Unique ℓ_2). There exists $\varepsilon > 0$ such that it is NP-hard to approximate uSVP in the ℓ_2 norm within a ratio $1 + \varepsilon$, even under randomized, sub-exponential-time reductions.

Acknowledgments

We thank Dor Minzer, Itamar Rot, Beata Kubis and Yonatan Pogrebinsky for helpful discussions and comments.

References

- [AC88] Noga Alon and Fan RK Chung. Explicit construction of linear sized tolerant networks. *Discrete Mathematics*, 72(1-3):15–19, 1988.
- [AD97] Miklós Ajtai and Cynthia Dwork. A public-key cryptosystem with worst-case/average-case equivalence. In *Proceedings of the twenty-ninth annual ACM symposium on Theory of computing*, pages 284–293, 1997.
- [AD07] Miklós Ajtai and Cynthia Dwork. The first and fourth public-key cryptosystems with worst-case/average-case equivalence. In *Electronic colloquium on computational complexity (ECCC)*, volume 14. Citeseer Princeton, NJ, USA, 2007.
- [AD16] Divesh Aggarwal and Chandan Dubey. Improved hardness results for unique shortest vector problem. *Information Processing Letters*, 116(10):631–637, 2016.
- [AFG13] Martin R Albrecht, Robert Fitzpatrick, and Florian Göpfert. On the efficacy of solving lwe by reduction to unique-svp. In *International Conference on Information Security and Cryptology*, pages 293–310. Springer, 2013.
- [AGVW17] Martin R Albrecht, Florian Göpfert, Fernando Virdia, and Thomas Wunderer. Revisiting the expected cost of solving usvp and applications to lwe. In *International Conference on the Theory and Application of Cryptology and Information Security*, pages 297–322. Springer, 2017.
- [Ajt98] Miklós Ajtai. The shortest vector problem in ℓ_2 is np-hard for randomized reductions. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 10–19, 1998.
- [AR05] Dorit Aharonov and Oded Regev. Lattice problems in $\text{np} \cap \text{conp}$. *J. ACM*, 52(5):749–765, September 2005.
- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, 1998.
- [BGPSD23] Huck Bennett, Atul Ganju, Pura Peetathawatchai, and Noah Stephens-Davidowitz. Just how hard are rotations of $\mathbb{Z}^n$? algorithms and cryptography with the simplest lattice. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 252–281. Springer, 2023.

- [BP20] Huck Bennett and Chris Peikert. Hardness of Bounded Distance Decoding on Lattices in ℓ_p Norms. In Shubhangi Saraf, editor, *35th Computational Complexity Conference (CCC 2020)*, volume 169 of *Leibniz International Proceedings in Informatics (LIPIcs)*, pages 36:1–36:21, Dagstuhl, Germany, 2020. Schloss Dagstuhl – Leibniz-Zentrum für Informatik.
- [BP23] Huck Bennett and Chris Peikert. Hardness of the (approximate) shortest vector problem: A simple proof via reed-solomon codes. In Nicole Megow and Adam D. Smith, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2023, September 11-13, 2023, Atlanta, Georgia, USA*, volume 275 of *LIPIcs*, pages 37:1–37:20. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2023.
- [BSW16] Shi Bai, Damien Stehlé, and Weiqiang Wen. Improved reduction from the bounded distance decoding problem to the unique shortest vector problem in lattices. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPIcs*, pages 76:1–76:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [Cai98] Jin-Yi Cai. A relation of primal-dual lattices and the complexity of shortest lattice vector problem. *Theoretical Computer Science*, 207(1):105–116, 1998.
- [CN11] Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology - ASIACRYPT 2011 - 17th International Conference on the Theory and Application of Cryptology and Information Security, Seoul, South Korea, December 4-8, 2011. Proceedings*, volume 7073 of *Lecture Notes in Computer Science*, pages 1–20. Springer, 2011.
- [DFK⁺99] Irit Dinur, Eldar Fischer, Guy Kindler, Ran Raz, and Shmuel Safra. PCP characterizations of NP: towards a polynomially-small error-probability. In Jeffrey Scott Vitter, Lawrence L. Larmore, and Frank Thomson Leighton, editors, *Proceedings of the Thirty-First Annual ACM Symposium on Theory of Computing, May 1-4, 1999, Atlanta, Georgia, USA*, pages 29–40. ACM, 1999.
- [DFK⁺11] Irit Dinur, Eldar Fischer, Guy Kindler, Ran Raz, and Shmuel Safra. PCP characterizations of NP: toward a polynomially-small error-probability. *Comput. Complex.*, 20(3):413–504, 2011.
- [Din02] Irit Dinur. Approximating svp_∞ to within almost-polynomial factors is np-hard. *Theoretical Computer Science*, 285(1):55–71, 2002.
- [DKK⁺17] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. *Electron. Colloquium Comput. Complex.*, TR17-094, 2017.
- [DKK⁺18] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In Ilias Diakonikolas, David Kempe, and Monika Henzinger, editors, *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 376–389. ACM, 2018.
- [DKRS03] I. Dinur, Guy Kindler, R. Raz, and S. Safra. Approximating cvp to within almost-polynomial factors is np-hard. *Combinatorica*, 23:205–243, 04 2003.
- [DKS98] I. Dinur, G. Kindler, and S. Safra. Approximating CVP to within almost-polynomial factors is NP-hard. In *Proc. 39th IEEE Symposium on Foundations of Computer Science*, 1998.
- [GG98] Oded Goldreich and Shafi Goldwasser. On the limits of non-approximability of lattice problems. In *Proceedings of the thirtieth annual ACM symposium on Theory of computing*, pages 1–9, 1998.

- [GGH96] Oded Goldreich, Shafi Goldwasser, and Shai Halevi. Public-key cryptosystems from lattice reduction problems. *Electron. Colloquium Comput. Complex.*, TR96-056, 1996.
- [HLW06] Shlomo Hoory, Nathan Linial, and Avi Wigderson. Expander graphs and their applications. *Bulletin of the American Mathematical Society*, 43(4):439–561, 2006.
- [HPS98] Jeffrey Hoffstein, Jill Pipher, and Joseph H. Silverman. NTRU: A ring-based public key cryptosystem. In Joe Buhler, editor, *Algorithmic Number Theory, Third International Symposium, ANTS-III, Portland, Oregon, USA, June 21-25, 1998, Proceedings*, volume 1423 of *Lecture Notes in Computer Science*, pages 267–288. Springer, 1998.
- [HR12] Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *Theory Comput.*, 8(1):513–531, 2012.
- [HR18] Ishay Haviv and Oded Regev. Tensor-based hardness of the shortest vector problem to within almost polynomial factors. *arXiv preprint arXiv:1806.04087*, 2018.
- [JX24] Baolong JIN and Rui XUE. Fine-grained hardness of the unique shortest vector problem in lattices. *SCIENTIA SINICA Informationis*, 54(12):2727, 2024.
- [Kan87] R. Kannan. Minkowski’s convex body theorem and integer programming. *Mathematics of Operations Research*, 12:415–440, 1987.
- [Kho03] S. Khot. Hardness of approximating the shortest vector problem in high L_p norms. In *Proc. 44th IEEE Symposium on Foundations of Computer Science*, 2003.
- [Kho05] Subhash Khot. Hardness of approximating the shortest vector problem in lattices. *J. ACM*, 52(5):789–808, sep 2005.
- [KM22] Tali Kaufman and Dor Minzer. Improved optimal testing results from global hypercontractivity. In *63rd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2022, Denver, CO, USA, October 31 - November 3, 2022*, pages 98–109. IEEE, 2022.
- [KM25] Tali Kaufman and Dor Minzer. Improved optimal testing results from global hypercontractivity. *SIAM J. Comput.*, 54(3):625–663, 2025.
- [KMS17] Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and grassmann graphs. In Hamed Hatami, Pierre McKenzie, and Valerie King, editors, *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 576–589. ACM, 2017.
- [KMS23] Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in Grassmann graph have near-perfect expansion. *Annals of Mathematics*, 198(1):1 – 92, 2023.
- [KS01] S Ravi Kumar and D Sivakumar. On the unique shortest lattice vector problem. *Theoretical computer science*, 255(1-2):641–648, 2001.
- [KT08] Than Quang Khoat and Nguyen Hong Tan. Unique shortest vector problem for max norm is np-hard. *Cryptology ePrint Archive*, 2008.
- [LLM06] Yi-Kai Liu, Vadim Lyubashevsky, and Daniele Micciancio. On bounded distance decoding for general lattices. In Josep Díaz, Klaus Jansen, José D. P. Rolim, and Uri Zwick, editors, *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques*, pages 450–461, Berlin, Heidelberg, 2006. Springer Berlin Heidelberg.
- [LM09] Vadim Lyubashevsky and Daniele Micciancio. On bounded distance decoding, unique shortest vectors, and the minimum distance problem. In *Annual International Cryptology Conference*, pages 577–594. Springer, 2009.

- [LP11] Richard Lindner and Chris Peikert. Better key sizes (and attacks) for lwe-based encryption. In Aggelos Kiayias, editor, *Topics in Cryptology - CT-RSA 2011 - The Cryptographers' Track at the RSA Conference 2011, San Francisco, CA, USA, February 14-18, 2011. Proceedings*, volume 6558 of *Lecture Notes in Computer Science*, pages 319–339. Springer, 2011.
- [Mic98] D. Micciancio. *On the hardness of the shortest vector problem*. PhD Thesis, MIT, 1998.
- [Mic01] Daniele Micciancio. The shortest vector in a lattice is hard to approximate to within some constant. *SIAM journal on Computing*, 30(6):2008–2035, 2001.
- [Mic12] Daniele Micciancio. Inapproximability of the shortest vector problem: Toward a deterministic reduction. *Theory of Computing*, 8(1):487–512, 2012.
- [Min22] Dor Minzer. *On Monotonicity Testing and the 2-to-2 Games Conjecture*, volume 49 of *ACM Books*. ACM, 2022.
- [MR10] Dana Moshkovitz and Ran Raz. Sub-constant error probabilistically checkable proof of almost-linear size. *Comput. Complex.*, 19(3):367–422, 2010.
- [MZ23] Dor Minzer and Kai Zhe Zheng. Optimal testing of generalized reed-muller codes in fewer queries. In *64th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2023, Santa Cruz, CA, USA, November 6-9, 2023*, pages 206–233. IEEE, 2023.
- [MZ24] Dor Minzer and Kai Zhe Zheng. Near optimal alphabet-soundness tradeoff pcps. In Bojan Mohar, Igor Shinkar, and Ryan O'Donnell, editors, *Proceedings of the 56th Annual ACM Symposium on Theory of Computing, STOC 2024, Vancouver, BC, Canada, June 24-28, 2024*, pages 15–23. ACM, 2024.
- [Pei16] Chris Peikert. A decade of lattice cryptography. *Found. Trends Theor. Comput. Sci.*, 10(4):283–424, 2016.
- [Reg04] Oded Regev. New lattice-based cryptographic constructions. *Journal of the ACM (JACM)*, 51(6):899–942, 2004.
- [Reg09] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. *J. ACM*, 56(6), September 2009.
- [RS97] Ran Raz and Shmuel Safra. A sub-constant error-probability low-degree test, and a sub-constant error-probability PCP characterization of NP. In Frank Thomson Leighton and Peter W. Shor, editors, *Proceedings of the Twenty-Ninth Annual ACM Symposium on the Theory of Computing, El Paso, Texas, USA, May 4-6, 1997*, pages 475–484. ACM, 1997.
- [Saf22] Muli Shmuel Safra. Mathematics of computation through the lens of linear equations and lattices. In *Proc. Int. Cong. Math*, volume 6, pages 4914–4969, 2022.
- [SD15] Noah Stephens-Davidowitz. Search-to-decision reductions for lattice problems with approximation factors (slightly) greater than one. *arXiv preprint arXiv:1512.04138*, 2015.
- [vE81] Boas P van Emde. Another np-complete partition problem and the complexity of computing short vectors in lattices. *TR*, 1981.
- [VV85] Leslie G Valiant and Vijay V Vazirani. Np is as easy as detecting unique solutions. In *Proceedings of the seventeenth annual ACM symposium on Theory of computing*, pages 458–463, 1985.

8 Appendix A

To make the paper self-contained, we follow an analysis of Kaufman and Minzer [KM22], proving Fact 2.11. We assume basic familiarity with eigenvalues, characters, and Cayley graphs. If needed, the survey of Hoory, Linial, and Wigderson [HLW06] contains all the necessary background and far more.

The proof consists of the following steps. First, it is possible to move from $G_{\mathcal{PvP}}$ to a certain Cayley graph (Actually, as seen in [KM22], from a more general case known as the Affine Grassmann graph). The eigenvalues of that Cayley graph are easy to compute. Once the eigenvalues are computed, the well-known *Expander Mixing Lemma* implies Fact 2.11.

8.1 The Cayley graph

Starting from the $\mathcal{PvP}$ graph over $\mathbb{F}_q^t$, $t > 2$, the vertices of our Cayley graph are the triplets $(s, x_1, x_2) \in (\mathbb{F}_q^t)^3$. Each vertex (s, x_1, x_2) has the following edges, described via a randomized process:

1. Uniformly sample $y \in \mathbb{F}_q^t$ and $\alpha, \beta, \gamma \in \mathbb{F}_q$.
2. Move to $(s + \alpha y, x_1 + \beta y, x_2 + \gamma y)$.

Denote the transition/random-walk matrix by M^{Cay} . Observe that M^{Cay} is symmetric, as starting from (s, x_1, x_2) and using y, α, β, γ is equivalent to starting from $(s + \alpha y, x_1 + \beta y, x_2 + \gamma y)$ and using $-y, \alpha, \beta, \gamma$.

The connection to $G_{\mathcal{PvP}}$ arises from associating each (s, x_1, x_2) with an affine subspace $s + \text{span}(x_1, x_2)$. Given a set $S \subseteq \mathbf{PL}(\mathbb{F}_q^t)$, denote:

$$S^* \stackrel{\text{def}}{=} \{(s, x_1, x_2) \in (\mathbb{F}_q^t)^3 \mid s + \text{span}(x_1, x_2) \in S\}$$

It turns out that the expansion of S and S^* are closely related. Thus, since the eigenvalues of Cayley graphs can be more easily calculated, in the subsequent sections we prove expansion in the Cayley graph. The following claim allows us to derive almost the same results for $G_{\mathcal{PvP}}$.

Claim 8.1. $\Phi(S) \geq \Phi(S^*) - \frac{1}{q} - \frac{1}{q^2}$

Proof. Suppose $(s, x_1, x_2) \in S^*$ and sample a random neighbor $(s + \alpha y, x_1 + \beta y, x_2 + \gamma y)$. There are few cases:

1. $\beta, \gamma = 0$, which happens with probability $\frac{1}{q^2}$.
2. $y \in \text{span}(x_1, x_2)$, with probability q^{2-t} . As $t = 3$, we receive $\frac{1}{q}$.
3. $\dim(\text{span}(x_1 + \beta y, x_2 + \gamma y)) < 2$. This case is contained in $y \in \text{span}(x_1, x_2)$, as a linear combination $c_1(x_1 + \beta y) + c_2(x_2 + \gamma y)$ could be rewritten as $c_1 \cdot x_1 + c_2 \cdot x_2 = (\beta c_1 + \gamma c_2)y$.
4. Otherwise, $(s + \alpha y) + \text{span}(x_1 + \beta y, x_2 + \gamma y)$ is an affine plane intersecting with $s + \text{span}(x_1, x_2)$ on a line. Moreover, it is distributed uniformly across such planes. The proof is elementary but slightly technical, and is written in the next paragraphs.

First, let us calculate the intersection. As $\beta \neq 0 \vee \gamma \neq 0$, either

$$\begin{aligned} s + \alpha y + \alpha \cdot \beta^{-1}(x_1 + \beta y) &= s + \alpha \cdot \beta^{-1}x_1 \in s + \text{span}(x_1, x_2) \\ \text{or } s + \alpha y + \alpha \cdot \gamma^{-1}(x_2 + \gamma y) &= s + \alpha \cdot \gamma^{-1}x_2 \in \text{span}(x_1, x_2) \end{aligned}$$

In addition, as $y \notin \text{span}(x_1, x_2)$, the planes differ. It's not hard to see that the intersection is a line with a gradient of:

$$\gamma(x_1 + \beta y) - \beta(x_2 + \gamma y) = \gamma x_1 - \beta x_2$$

Fixing β, γ , suppose w.l.o.g $\beta \neq 0$. The point $s + \alpha \cdot \beta^{-1}x_1$ is inside the intersection, so the line is:

$$\ell(t) = s + \alpha \cdot \beta^{-1}x_1 + t \cdot (\gamma x_1 - \beta x_2)$$

For different values of α , the lines are parallel. Thus, when sampling α, β, γ (independent of y), the intersection distributes uniformly over the lines. Sampling α, β, γ and then y , has the same distribution as uniformly sampling a line in $s + \text{span}(x_1, x_2)$, and then adding a random third point outside the plane.

To conclude the proof of [Claim 8.1](#), the first three cases occur with a probability of at most $\frac{1}{q} + \frac{1}{q^2}$. $(s + \alpha y, x_1 + \beta y, x_2 + \gamma y) \in S^*$ is equivalent to $s + \alpha y + \text{span}(x_1 + \beta y, x_2 + \gamma y) \in S$, so the fourth case corresponds to a random walk in the $G_{\mathcal{P}v\mathcal{P}}$ graph. It implies that:

$$\frac{1}{q} + \frac{1}{q^2} + (1 - \frac{1}{q} - \frac{1}{q^2})\Phi(S) \geq \Phi(S^*) \Rightarrow \Phi(S) \geq \Phi(S^*) - \frac{1}{q} - \frac{1}{q^2}$$

□

8.2 Bounding the Eigenvalues

Our next step is calculating the eigenvalues of M^{Cay} . Again, this was done before in [\[KM22\]](#). It is folklore that, for Cayley graphs, the characters are the eigenvectors of the graph. While the graph is weighted, this is still true. For a character χ_x , denote the corresponding eigenvalue as λ_x .

Claim 8.2. *For all $\vec{0} \neq x = (s, x_1, x_2) \in (\mathbb{F}_q^t)^3$, the eigenvalue λ_x is bounded by $|\lambda_x| \leq \frac{1}{q}$.*

Proof. A calculation of $M^{\text{Cay}}\chi_x$ in any coordinate shows:

$$\lambda_x = \sum_{u \in (\mathbb{F}_q^t)^3} M_{\vec{0}, u}^{\text{Cay}} \chi_x(u) = \mathbb{E}_{y, \alpha, \beta, \gamma} [\chi_x(\alpha y, \beta y, \gamma y)]$$

Reordering the term inside, we receive the following:

$$\mathbb{E}_{y, \alpha, \beta, \gamma} [\chi_x(\alpha y, \beta y, \gamma y)] = \mathbb{E}_{y, \alpha, \beta, \gamma} [\chi_{\alpha s + \beta x_1 + \gamma x_2}(y)] = \begin{cases} 1 & \alpha s + \beta x_1 + \gamma x_2 = 0 \\ 0 & \text{else} \end{cases}$$

The first case happens with a probability of $q^{-\dim \text{span}(s, x_1, x_2)}$. We assumed $x \neq \vec{0}$ so $|\lambda_x| \leq \frac{1}{q}$. □

8.3 Concluding Fact 2.11

One of the most fundamental ideas in the theory of expanders is that good expanders “look random”. The *Expander Mixing lemma* [\[AC88\]](#) states that for every d -regular graph $G = (V, E)$ and a subset $S \subseteq V$, if G has a good spectral expansion, the number of edges with endpoints in S and V (crossing the partition), is roughly $d \cdot |S| \frac{|V| - |S|}{|V|}$. Equivalently, $\Phi(S) \approx 1 - \frac{|S|}{|V|}$. While our Cayley graph is weighted, the classical proof still works.

Lemma 8.3 (Expander mixing lemma — weighted). *Let G be a weighted graph on vertices $[n]$ with a symmetric random-walk matrix $W \in [0, 1]^{n \times n}$. Denote the eigenvalues $\lambda_1 \geq \dots \geq \lambda_n$. For $\lambda > 0$, assume $\max_{i \neq 1} |\lambda_i| \leq \lambda$. For every $S \subset [n]$:*

$$|\Phi(S) - 1 + \frac{|S|}{n}| \leq \lambda$$

Proving [Fact 2.11](#) is immediate. [Lemma 8.3](#) states that for every $S \subset V_{\mathcal{P}v\mathcal{P}}$ in the $\mathcal{P}v\mathcal{P}$ graph and a corresponding $S^* \subseteq (\mathbb{F}_q^t)^3$ in the Cayley graph:

$$\Phi(S) + \frac{1}{q} + \frac{1}{q^2} \geq \Phi(S^*) \geq 1 - \frac{|S^*|}{q^{3t}} - \frac{1}{q} \Rightarrow \Phi(S) \geq 1 - \frac{2}{q} - \frac{1}{q^2} - \frac{|S^*|}{q^{3t}}$$

Observe that:

$$\frac{|S|}{|V_{\mathcal{P}v\mathcal{P}}|} = \frac{|S^*|}{|\{(s, x_1, x_2) \in (\mathbb{F}_q^t)^3 \mid \dim \text{span}(x_1, x_2) = 2\}|} \geq \frac{|S^*|}{q^{3t}}$$

So a loose bound on the expansion is ([Fact 2.11](#)):

$$\Phi(S) \geq 1 - \frac{|S|}{|V_{\mathcal{P}v\mathcal{P}}|} - \frac{3}{q}.$$