Near Optimal Alphabet-Soundness Tradeoff PCPs

Dor Minzer*
Massachusetts Institute of Technology
Cambridge, MA, USA

Kai Zhe Zheng†
Massachusetts Institute of Technology
Cambridge, MA, USA

ABSTRACT

We show that for all $\varepsilon > 0$, for sufficiently large prime power $q \in \mathbb{N}$, for all $\delta > 0$, it is NP-hard to distinguish whether a 2-Prover-1-Round projection game with alphabet size q has value at least $1-\delta$, or value at most $1/q^{1-\varepsilon}$. This establishes a nearly optimal alphabet-to-soundness tradeoff for 2-query PCPs with alphabet size q, improving upon a result of [Chan 2016]. Our result has the following implications:

(1) Near optimal hardness for Quadratic Programming: it is NP-hard to approximate the value of a given Boolean Quadratic Program within factor $(\log n)^{1-o(1)}$ under quasi-polynomial time reductions. This result improves a result of [Khot-Safra 2013] and nearly matches the performance of the best known approximation algorithm [Megrestki 2001, Nemirovski-Roos-Terlaky 1999 Charikar-Wirth 2004] that achieves a factor of $O(\log n)$.
(2) Bounded degree 2-CSP's: under randomized reductions, for sufficiently large $d > 0$, it is NP-hard to approximate the value of 2-CSPs in which each variable appears in at most d constraints within factor $(1-o(1))\frac{d}{2}$, improving upon a recent result of [Lee-Manurangsi 2023].
(3) Improved hardness results for connectivity problems: using results of [Laekhanukit 2014] and [Manurangsi 2019], we deduce improved hardness results for the Rooted k-Connectivity Problem, the Vertex-Connectivity Survivable Network Design Problem and the Vertex-Connectivity k-Route Cut Problem.

CCS CONCEPTS

• **Theory of computation → Computational complexity and cryptography**.

KEYWORDS

Probabilistically Checkable Proofs, Hardness of Approximation, Label Cover, 2-Prover-1-Round Games

ACM Reference Format:
Dor Minzer and Kai Zhe Zheng. 2024. Near Optimal Alphabet-Soundness Tradeoff PCPs. In *Proceedings of the 56th Annual ACM Symposium on Theory of Computing (STOC '24), June 24–28, 2024, Vancouver, BC, Canada.* ACM, New York, NY, USA, 9 pages. https://doi.org/10.1145/3618260.3649606

*Supported by NSF CCF award 2227876 and NSF CAREER award 2239160.
†Supported by the NSF GRFP DGE-2141064.

STOC '24, June 24–28, 2024, Vancouver, BC, Canada

ACM ISBN 979-8-4007-0383-6/24/06
https://doi.org/10.1145/3618260.3649606

1 INTRODUCTION

The PCP theorem is a fundamental result in theoretical computer science with many equivalent formulations [2, 3, 18]. One of the formulations asserts that there exists $\varepsilon > 0$ such that given a satisfiable 3-SAT formula ϕ, it is NP-hard to find an assignment that satisfies at least $(1-\varepsilon)$ fraction of the constraints. The PCP theorem has a myriad of applications within theoretical computer science, and of particular interest to this paper are applications of PCP to hardness of approximation.

The vast majority of hardness of approximation result are proved via reductions from the PCP theorem above. Oftentimes, to get a strong hardness of approximation result, one must first amplify the basic PCP theorem above into a result with stronger parameters [17, 21, 22, 26] (see [41] for a survey). To discuss these parameters, it is often convenient to view the PCP through the problem of 2-Prover-1-Round Games, which we define next.[1]

Definition 1.1. *An instance Ψ of 2-Prover-1-Round Games consists of a bipartite graph $G = (L \cup R, E)$, alphabets Σ_L and Σ_R and a collection of constraints $\Phi = \{\phi_e\}_{e\in E}$, which for each edge $e \in E$ specifies a constraint map $\phi_e : \Sigma_L \to \Sigma_R$.*

(1) *The alphabet size of Ψ is defined to be $|\Sigma_L| + |\Sigma_R|$.*
(2) *The value of Ψ is defined to be the maximum fraction of edges $e \in E$ that can be satisfied by any assignment. That is,*

$$\mathrm{val}(\Psi) = \max_{\substack{A_L : L\to\Sigma_L \\ A_R : R\to\Sigma_R}} \frac{|\{e=(u,v)\in E \mid \phi_e(A_L(u)) = A_R(v)\}|}{|E|}.$$

The combinatorial view of 2-Prover-1-Round Games has its origins in an equivalent, active view in terms of a game between a verifier and two all powerful provers, which is sometimes more intuitive. The verifier and the two provers have access to an instance Ψ of 2-Prover-1-Round Games, and the provers may agree beforehand on a strategy; after this period they are not allowed to communicate. The verifier then picks a random edge, $e = (u, v)$, from the 2-Prover-1-Round game, sends u to the first prover, sends v to the second prover, receives a label in response from each one of them, and finally checks that the labels satisfy the constraint ϕ_e. If so, then the verifier accepts. It is easy to see that the value of the 2-Prover-1-Round game is equal to the acceptance probability of the verifier under the best strategy of the provers.

In the language of 2-Prover-1-Round Games, the majority of hardness of approximation results are proved by combining the basic PCP theorem [2, 3, 18] with Raz's parallel repetition theorem [38], which together imply the following result:

[1] Strictly speaking, the notion below is referred to in the literature as projection 2-Prover-1-Round games. We omit the more general definition as we do not discuss non-projection games in this paper.

Theorem 1.2. *There exists $\gamma > 0$ such that for sufficiently large R, given a 2-Prover-1-Round game Ψ with alphabet size R, it is NP-hard to distinguish between the following two cases:*

(1) *YES case:* $\mathrm{val}(\Psi) = 1$.
(2) *NO case:* $\mathrm{val}(\Psi) \leqslant \frac{1}{R^{\gamma}}$.

For many applications, one only requires that the soundness error of the PCP is small. Namely, that $\mathrm{val}(\Psi)$ is arbitrarily small in the "NO case". For certain applications however, more is required: not only must the soundness error be small – but it must also be small in terms of the alphabet size. The tradeoff between the soundness error of the PCP and the alphabet size of the PCP is the main focus of this paper.

With respect to this tradeoff, it is clear that the best result one may hope for in Theorem 1.2 is $\gamma = 1 - o(1)$ since a random assignment to Ψ satisfies, in expectation, at least $\frac{1}{R}$ fraction of the constraints. In terms of results, combining the PCP theorem with Raz's parallel repetition theorem gives $\gamma > 0$ that is an absolute, but tiny constant. Towards a stronger tradeoff, Khot and Safra [27] showed that Theorem 1.2 holds for $\gamma = 1/6$ with imperfect completeness (i.e., $\mathrm{val}(\Psi) \geqslant 1 - o(1)$ instead of $\mathrm{val}(\Psi) = 1$ in the "YES case"). The result of Khot and Safra was improved by Chan [8], who showed (using a completely different set of techniques) that Theorem 1.2 holds for $\gamma = 1/2 - o(1)$, again with imperfect completeness.

In the remainder of this paper we will describe our main results and give an overview of our PCP construction. Additional details, including proofs, can be found in the full version of this paper [33].

1.1 Main Results

In this section we explain the main results of this paper.

1.1.1 Near Optimal Alphabet vs Soundness Tradeoff. The main result of this work improves upon all prior results, and shows that one may take $\gamma = 1 - o(1)$ in Theorem 1.2, again with imperfect completeness. Formally, we show:

Theorem 1.3. *For all $\varepsilon, \delta > 0$, for sufficiently large R, given a 2-Prover-1-Round game Ψ, it is NP-hard to distinguish between the following two cases:*

(1) *YES case:* $\mathrm{val}(\Psi) \geqslant 1 - \delta$.
(2) *NO case:* $\mathrm{val}(\Psi) \leqslant \frac{1}{R^{1-\varepsilon}}$.

Theorem 1.3 gives a near optimal tradeoff between the alphabet size of a PCP and the soundness of a PCP, improving upon the result of Chan [8]. Moreover, Theorem 1.3 has several applications to combinatorial optimization problems, which we discuss below. We remark that most of these applications require additional features from the instances produced in Theorem 1.3 which we omit from its formulation for the sake of clarity. For instance, one application requires a good tradeoff between the size of the instance and the size of the alphabet, which our construction achieves (see the discussion following Theorem 1.4). Other applications require the underlying constraint graph to be bounded-degree bi-regular graph, which our construction also achieves, after mild modifications detailed in [29].

1.1.2 Application: NP-Hardness of Approximating Quadratic Programs. Theorem 1.3 has an application to the hardness of approximating the value of Boolean Quadratic Programming, as we explain next.

An instance of Quadratic programming consists of a quadratic form $Q(x) = \sum_{i,j=1}^{n} a_{i,j} x_i x_j$ where $a_{i,i} = 0$ for all i, and one wishes to maximize $Q(x)$ over $x \in \{-1, 1\}^n$. This problem is known to have an $O(\log n)$ approximation algorithm [9, 31, 35], and is known to be quasi-NP-hard to approximate within factor $(\log n)^{1/6-o(1)}$ [1, 27]. That is, unless NP has a quasi-polynomial time algorithm, no polynomial time algorithm can approximate Quadratic Programming to within factor $(\log n)^{1/6-o(1)}$. As a first application of Theorem 1.3, we improve the hardness result of Khot and Safra:

Theorem 1.4. *It is quasi-NP-hard to approximate Quadratic Programming to within a factor of $(\log n)^{1-o(1)}$.*

Theorem 1.4 is proved via a connection between 2-Prover-1-Round Games and Quadratic Programming due to Arora, Berger, Hazan, Kindler, and Safra [1]. This connections requires a good tradeoff between the alphabet size, the soundness error, and the size of the PCP. Fortunately, the construction in Theorem 1.4 has a sufficiently good tradeoff between all of these parameters: letting N be the size of the instance, the alphabet size can be taken to be $(\log N)^{1-o(1)}$ and the soundness error can be taken to be $(\log N)^{-1+o(1)}$. [2]

Relevance to the sliding scale conjecture: It is worth noting that using our techniques, we do not know how to achieve soundness error that is smaller than inversely poly-logarithmic in the instance size. As such, our techniques have no bearing on the sliding scale conjecture, which is concerned with getting soundness error that is inversely polynomial in the instance size. This seems to be a bottleneck of any PCP construction that is based on the covering property. In fact, assuming ETH, any quasi-polynomial PCP construction achieving soundness error, say, $1/(\log N)^2$ would necessarily need to have almost polynomial alphabet size (since the reduction to Quadratic Solvability would give an algorithm that runs roughly in time exponential in the alphabet size), which is the opposite of what our techniques give. With this in mind, we would like to mention a closely related, recent conjecture made in [10], which is a sort of a mixture between d-to-1 games and the sliding scale conjecture. This conjecture is motivated by improved hardness results for densest k-subgraph style problems, and focuses on the relation between the instance size and the soundness error (allowing the alphabet to be quite large). It may be possible that the ideas from the current paper can help make progress towards this conjecture.

1.1.3 Application: NP-hardness of Approximating Bounded Degree 2-CSPs. Theorem 1.3 has an application to the hardness of approximating the value of 2-CSPs with bounded degree, as we explain next.

[2] We remark that the result of Chan [8] does not achieve a good enough trade-off between the alphabet size and the instance size due to the use of the long-code, and therefore it does not yield a strong inapproximability result for Quadratic Programming.

An instance Ψ of 2-CSP, say $\Psi = (X, C, \Sigma)$, consists of a set of variables X, a set of constraints C and an alphabet Σ. Each constraint in C has the form $P(x_i, x_j) = 1$ where $P\colon \Sigma \times \Sigma \to \{0,1\}$ is a predicate (which may be different in distinct constraints). The degree of the instance Ψ is defined to be the maximum, over variables $x \in X$, of the number of constraints that x appears in. The goal is to find an assignment $A\colon X \to \Sigma$ that satisfies as many of the constraints as possible.

There is a simple $\frac{d+1}{2}$ approximation algorithm for the 2-CSP problem for instances with degree at most d. Lee and Manurangsi proved a nearly matching $\left(\frac{1}{2} - o(1)\right) d$ hardness of approximation result assuming the Unique-Games Conjecture [29]. Unconditionally, they show the problem to be NP-hard to approximate within factor $\left(\frac{1}{3} - o(1)\right) d$ under randomized reductions.

Using the ideas of Lee and Manurangsi, our main result implies a nearly matching NP-hardness result for bounded degree 2-CSPs:

Theorem 1.5. *For all $\eta > 0$, for sufficiently large d, approximating the value of 2-CSPs with degree at most d within factor $\left(\frac{1}{2} - \eta\right) d$ is NP-hard under randomized reductions.*

As in [29], Theorem 1.5 has a further application to finding independent sets in claw free graphs. A k-claw $K_{1,k}$ is the $(k+1)$ vertex graph with a center vertex which is connected to all other k-vertices and has no other edges; a graph G is said to be k-claw free if G does not contain an induced k-claw graph. There is a polynomial time approximation algorithm for approximating the size of the largest independent set in a given k-claw free graph G within factor $\frac{k}{2}$ [4, 40], and a quasi-polynomial time approximation algorithm within factor $\left(\frac{1}{3} + o(1)\right) k$ [11]. As in [29], using ideas from [14] Theorem 1.5 implies that for all $\varepsilon > 0$, for sufficiently large k, it is NP-hard (under randomized reductions) to approximate the size of the largest independent set in a given k-claw free graph within factor $\left(\frac{1}{4} + \eta\right) k$. This improves upon the result of [29] who showed that the same result holds assuming the Unique-Games Conjecture.

1.1.4 Application: NP-hardness of Approximating Connectivity Problems. Using ideas of Laekhanukit [28] and the improvements by Manurangsi [30], Theorem 1.3 implies improved hardness of approximation results for several graph connectivitiy problems. More specifically, Theorem 1.3 combined with the results of [30] implies improvements to each one of the results outlined in table 1 in [28] by a factor of 2 in the exponent - with the exception of Rooted-k-Connectivity on directed graphs where a factor of 2 improvement is already implied by [30]. We briefly discuss the Rooted k-Connectivity Problem, but defer the reader to [28] for a detailed discussion of the remaining graph connectivity problems.

In the Rooted k-Connectivity problem there is a graph $G = (V, E)$, edge costs $c\colon E \to \mathbb{R}$, a root vertex $r \in V$ and a set of terminals $T \subseteq V \setminus \{r\}$. The goal is to find a sub-graph G' of smallest cost that for each $t \in T$, has at least k vertex disjoint paths from r to t. The problem admits $|T|$ trivial approximation algorithm (by applying minimum cost k-flow algorithm for each vertex in T), as well as an $O(k \log k)$ approximation algorithm [36].

Using the ideas of [28], Theorem 1.3 implies the following improved hardness of approximation results:

Theorem 1.6. *For all $\varepsilon > 0$, for sufficiently large k it is NP-hard to approximate the Rooted-k-Connectivity problem on undirected graphs to within a factor of $k^{1/5-\varepsilon}$, the Vertex-Connectivity Survivable Network Design Problem with connectivity parameters at most k to within a factor of $k^{1/3-\varepsilon}$, and the Vertex-Connectivity k-Route Cut Problem to within a factor of $k^{1/3-\varepsilon}$.*

We remark that in [7], a weaker form of hardness for the Vertex-Connectivity Survivable Network problem is proved. More precisely, they show an $\Omega(k^{1/3}/\log k)$ integrality gap for the set-pair relaxation of the problem. Our hardness result of $k^{1/3-\varepsilon}$ improves upon it, showing that (unless P=NP) no relaxation can yield a better than $k^{1/3-\varepsilon}$ factor approximation algorithm.

2 PRELIMINARIES

In this section we will describe some preliminary definitions and results. We first present the Grassmann graph and some Fourier analytic tools that are used in our analysis. We then state some hardness results regarding 3Linwhich will form the starting point of our PCP construction.

2.1 The Grassmann Graph

Throughout this section, we fix parameters n, ℓ with $1 \ll \ell \ll n$, and a prime power q. The Grassmann graph $\mathrm{Grass}_q(n, \ell)$ is defined as follows.

- The vertex set corresponds to the set of ℓ-dimensional subspaces $L \subseteq \mathbb{F}_q^n$.
- The edge set corresponds to all pairs (L, L') of ℓ-dimensional subspaces $L, L' \subseteq \mathbb{F}_q^n$ such that $\dim(L \cap L') = \ell - 1$.

We also write $\mathrm{Grass}_q(V, \ell)$ to denote the Grassmann graph on ℓ-dimensional subspaces V, where V is some large linear subspace. Finally, we denote by $L_2(\mathrm{Grass}_q(n, \ell))$ the set of complex valued functions $F\colon \begin{bmatrix} \mathbb{F}_q^n \\ \ell \end{bmatrix}_q \to \mathbb{C}$, where $\begin{bmatrix} \mathbb{F}_q^n \\ \ell \end{bmatrix}_q$ is the set of ℓ-dimensional subspaces in $\mathbb{F}_q^n$.

Zoom ins and Zoom outs. A feature of the Grassmann graph is that it contains many copies of lower dimensional Grassmann graphs as induced subgraphs. These subgraphs are precisely the zoom-ins and and zoom-outs referred to in the introduction, and they play a large part in the analysis of our inner PCP and final PCP. For subspaces $Q \subseteq W \subseteq \mathbb{F}_q^n$, let

$$\mathrm{Zoom}[Q, W] = \{L \in \mathrm{Grass}_q(n, \ell) \mid Q \subseteq L \subseteq W\}.$$

We refer to Q as a zoom-in and W as a zoom-out. When $W = \mathbb{F}_q^n$, $\mathrm{Zoom}[Q, W]$ is the zoom-in on Q, and when $Q = \{0\}$, $\mathrm{Zoom}[Q, W]$ is the zoom-out on W.

2.1.1 Pseudo-randomness over the Grassmann graph. One notion that will be important to us is (r, ε)-pseudo-randomness, which measures how much F can deviate from its expectation on a zoom-in/zoom-out restrictions of "size r". For all of our applications, F and G will both be indicator functions of some sets of vertices, so it will be helpful to think of this case for the remainder of the section. [3] Let

[3] We remark that the results we state have more general versions that apply to wider classes of functions. We refrain from stating them in this generality for sake of simplicity.

$\mu(F) = \mathbb{E}_{L\in \mathrm{Grass}_q(n,\ell)}[F(L)]$ (for indicator functions, this is simply the measure of the indicated set). For subspaces $Q \subseteq W \subseteq \mathbb{F}_q^n$, define

$$\mu_{Q,W}(F) = \mathop{\mathbb{E}}_{L\in \mathrm{Grass}_q(n,\ell)}[F(L) \mid Q \subseteq L \subseteq W].$$

Definition 2.1. *We say that a Boolean function $F\colon G(n,\ell) \to \{0,1\}$ is (r,ε)-pseudo-random if for all $Q \subseteq W \subseteq \mathbb{F}_q^n$ satisfying $\dim(Q) + \mathrm{codim}(W) = r$, we have*

$$\mu_{Q,W}(F) \leqslant \varepsilon.$$

We will often say that a set $S \subseteq \mathrm{Grass}_q(n,\ell)$ is (r,ε)-pseudo-random if its indicator function is. Because the Grassmann graph is not a small-set expanders, there are small sets in it that do not look "random" with respect to some combinatorial counting measures (such as edges between sets, expansion and so on). Intuitively, a small set S which is highly pseudo-random will exhibit random-like structure with respect to several combinatorial measures of interest, and the two lemmas below are instantiations of it required in our proof. The proof proceed by reducing them to similar statements about the Bi-linear scheme, which can then be proved directed by appealing to global hypercontractivity results of [15, 16].

For the analysis of the inner PCP, we require the following lemma, which bounds the number of edges between a subset, $\mathcal{L}$, of $\mathrm{Grass}_q(n,2\ell)$, and $\mathrm{Grass}_q(n,2(1-\delta)\ell)$ when $\mathcal{L}$ is (r,ε)-pseudo-random.

Lemma 2.2. *Let $F\colon \mathrm{Grass}_q(n,2\ell) \to \{0,1\}$ and $G\colon \mathrm{Grass}_q(n,2(1-\delta)\ell) \to \{0,1\}$ be Boolean functions such that*

$$\mathop{\mathbb{E}}_{L}[F(L)] = \alpha, \qquad \mathop{\mathbb{E}}_{R} G(R)] = \beta,$$

and suppose that F is (r,ε) pseudo-random. Then for all $t \geqslant 4$ that are powers of 2,

$$\langle \mathcal{T}F, G\rangle \leqslant q^{O_{t,r}(1)}\beta^{(t-1)/t}\varepsilon^{2t/(2t-1)} + q^{-r\delta\ell}\sqrt{\alpha\beta}.$$

2.2 Hardness of 3LIN

In this section we cite several hardness of approximation results for the problem of solving linear equations over finite fields, which are the starting point of our reduction. We begin by defining the 3Lin and the Gap3Lin problem.

Definition 2.3. *For a prime power q, an instance of* 3Lin *is (X, Eq) which consists of a set of variables X and a set of linear equations* Eq *over $\mathbb{F}_q$. Each equation in* Eq *depends on exactly three variables in X, each variable appears in at most* 10 *equations, and any two distinct equations in* Eq *share at most a single variable.*

The goal in the 3Lin problem is to find an assignment $A\colon X \to \mathbb{F}_q$ satisfying as many of the equations in E as possible. The maximum fraction of equations that can be satisfied is called the value of the instance. We remark that usually in the literature, the condition that two equations in E share at most a single variable is not included in the definition of 3Lin, as well the the bound on the number of occurences of each variable.

For $0 < s < c \leqslant 1$, the problem $\mathrm{Gap3Lin}[c,s]$ is the promise problem wherein the input is an instance (X,E) of 3Lin promised to either have value at least c or at most s, and the goal is to distinguish between these two cases. The problem $\mathrm{Gap3Lin}[c,s]$ with various settings of c and s will be the starting point for our reductions.

To prove Theorem 1.3, we shall use the classical result of Håstad [22]. This result says that for general 3Lin instances (i.e., without the additional condition that two equations share at most a single variable), the problem $\mathrm{Gap3Lin}[1-\varepsilon, 1/q+\varepsilon]$ is NP-hard for all constant $q \in \mathbb{N}$ and $\varepsilon > 0$. This result implies the following theorem by elementary reductions:

Theorem 2.1. *There exists $s < 1$ such that for every constant $\eta > 0$ and prime q, $\mathrm{Gap3Lin}\,[1-\eta, s]$ is* NP*-hard.*

To prove Theorem 1.4 we will need a hardness result for 3Lin with completeness close to 1, and we will use a hardness result of Khot and Ponnuswami [26]. Once again, their result does not immediately guarantee the fact that any two equations share at most a single variable, however once again this property may be achieved by an elementary reduction.

Theorem 2.2. *There is a reduction from* SAT *with size n to an instance of $\mathrm{Gap3Lin}[1-\eta, 1-\varepsilon]$ of size N over a field $\mathbb{F}_q$ of characteristic 2, where,*

- *Both N and the running time of the reduction are bounded by $2^{O(\log^2 n)}$*
- *$\eta \leqslant 2^{-\Omega(\sqrt{\log N})}$.*
- *$\varepsilon \geqslant \Omega\left(\frac{1}{\log^3 N}\right)$.*

3 THE PCP CONSTRUCTION

Theorem 1.3 is proved by composing an inner PCP and an outer PCP. Both of these components incorporate ideas from the proof of the 2-to-1 Games Theorem. The outer PCP is constructed using smooth parallel repetition [24, 27] while the inner PCP is based on the Grassmann graph [12, 13, 24, 25].

The novelty in this current paper, in terms of techniques, is twofold. First, we must consider a Grassmann test in a different regime of parameters (as otherwise we would not be able to get a good alphabet to soundness tradeoff) and in a regime of much lower soundness error. These differences complicate matters considerably. Second, our soundness analysis is more involved than that of the 2-to-1-Games Theorem. As is the case in [12, 13, 24, 25], we too use global hyperconractivity, but we do so more extensively. We also require quantitatively stronger versions of global hypercontractivity over the Grasssmann graph which are due to [16]. In addition, our analysis incorporates ideas from the plane versus plane test and direct product testing [23, 32, 39], from classical PCP theory [27], as well as from error correcting codes [19]. All of these tools are necessary to prove our main technical statement – Lemma 3.1 below – which is a combinatorial statement that may be of independent interest.

We now elaborate on each one of the components separately.

3.1 The Inner PCP

Our Inner PCP is based on the subspace vs subspace low degree test. Below, we first give a general overview of the objective in low-degree testing. We then discuss the traditional notion of soundness as well as a non-traditional notion of soundness for low-degree tests. Finally, we explain the low-degree test used in this paper, the

notion of soundness that we need from it, and the way that this notion of soundness is used.

Low degree tests in PCPs. Low degree tests have been have a vital component in PCPs since their inception, and much attention has been devoted to improving their various parameters. The goal in low-degree testing is to encode a low-degree function $f: \mathbb{F}_q^n \to \mathbb{F}_q$ via a table (or a few tables) of values, in a way that allows for local testing. Traditionally, one picks a parameter $\ell \in \mathbb{N}$ (which is thought of as a constant and is most often just 2) and encodes the function f by the table T of restrictions of f to ℓ-dimensional affine subspaces of $\mathbb{F}_q^n$. For the case $\ell = 2$, the test associated with this encoding is known as the Plane vs Plane test [39]. The Plane vs Plane test proceeds by picking two planes P_1, P_2 intersecting on a line, and then checking that $T[P_1]$ and $T[P_2]$ agree on $P_1 \cap P_2$. It is easy to see that the test has perfect completeness, namely that a valid table of restrictions T passes the test with probability 1. In the other direction, the soundness error of the test – which is a converse type statement – is much less clear (and is crucial towards applications in PCP). In the context of the Plane vs Plane test, it is know that if a table T, that assigns to each plane a degree d function, passes the Plane vs Plane test with probability $\varepsilon \geqslant q^{-c}$ (where $c > 0$ is a small absolute constant), then there is a degree d function f such that $T[P] \equiv f|_P$ on at least $\Omega(\varepsilon)$ fraction of the planes.

Nailing down the value of the constant c for which soundness holds is an interesting open problem which is related to soundness vs alphabet size vs instance size tradeoff in PCPs [5, 32, 34]. Currently, the best known analysis for the Plane vs Plane test [34] shows that one may take $c = 1/8$. Better analysis is known for higher dimensional encoding [5, 32], and for the 3-dimensional version of it a near optimal soundness result is known [32].

Low degree tests in this paper. In the context of the current paper, we wish to encode linear functions $f: \mathbb{F}_q^n \to \mathbb{F}_q$, and we do so by the subspaces encoding. Specifically, we set integer parameters $\ell_1 \geqslant \ell_2$, and encode the function f using the table T_1 of the restrictions of f to all ℓ_1-dimensional linear subspaces of $\mathbb{F}_q^n$, and the table T_2 of the restrictions of f to all ℓ_2-dimensional linear subspaces of $\mathbb{F}_q^n$. The test we consider is the natural inclusion test:

(1) Sample a random ℓ_1-dimensional subspace $L_1 \subseteq \mathbb{F}_q^n$ and a random ℓ_2-dimensional subspace $L_2 \subseteq L_1$.
(2) Read $T_1[L_1], T_2[L_2]$ and accept if they agree on L_2.

As is often the case, the completeness of the test – namely the fact that valid tables T_1, T_2 pass the test with probability 1 – is clear. The question of most interest then is with regards to the soundness of the test. Namely, what is the smallest ε such that any two tables T_1 and T_2 that assign linear functions to subspaces and pass the test with probability ε, must necessarily "come from" a legitimate linear function f?

Traditional notion of soundness. As the alphabet vs soundness tradeoff is key to the discussion herein, we begin by remarking that the alphabet size of the above encoding is $q^{\ell_1} + q^{\ell_2} = \Theta(q^{\ell_1})$ (since there are q^{ℓ} distinct linear functions on a linear space of dimension ℓ over $\mathbb{F}_q$). Thus, ideally we would like to show that the soundness error of the above test is $q^{-(1-o(1))\ell_1}$. Alas, this is false. Indeed, it turns out that one may construct assignments that pass the test with probability at least $\Omega(\max(q^{-\ell_2}, q^{\ell_2-\ell_1}))$ that do not have significant correlation with any linear function f:

(1) Taking T_1, T_2 randomly by assigning to each subspace a random linear function, one can easily see that the test passes with probability $\Theta(q^{-\ell_2})$.
(2) Taking linear subspaces $W_1, \ldots, W_{100q^{\ell_1}} \subseteq \mathbb{F}_q^n$ of co-dimension 1 randomly, and a random linear function $f_i: W_i \to \mathbb{F}_q$ for each i, one may choose T_1 and T_2 as follows. For each L_1, pick a random i such that $L_1 \subseteq W_i$ (if such i exists) and assign $T_1[L_1] = f_i|_{L_1}$. For each L_2, pick a random i such that $L_2 \subseteq W_i$ (if such i exists) and assign $T_2[L_2] = f_i|_{L_2}$. Taking $L_2 \subseteq L_1$ randomly, one sees that with constant probability L_2 has $\Theta(q^{\ell_1-\ell_2})$ many possible i's, L_1 has $\Theta(1)$ many possible i's and furthermore there is at least one i that is valid for both of them. With probability $\Omega(q^{\ell_2-\ell_1})$ this common i is chosen for both L_1 and L_2, and in this case, the test on (L_1, L_2) passes. It follows that, in expectation, T_1, T_2 pass the test with probability $\Omega(q^{\ell_2-\ell_1})$.

In light of the above, it makes sense that the best possible alphabet vs soundness tradeoff we may achieve with the subspace encoding is by taking $\ell_2 = \ell_1/2$. Such a setting of the parameters would give alphabet size $R = q^{\ell_1}$ and (possibly) soundness error $\Theta(1/\sqrt{R})$. There are several issues with this setting however. First, this tradeoff is not good enough for our purposes (which already rules out this setting of parameters). Second, we do not know how to prove that the soundness error of the test is $\Theta(1/\sqrt{R})$ (the best we can do is quadratically off and is $\Theta(1/R^{1/4})$). To address both of these issues, we must venture beyond the traditional notion of soundness.

Non-traditional notion of soundness. The above test was first considered in the context of the 2-to-1 Games Theorem, wherein one takes $q = 2$ and $\ell_2 = \ell_1 - 1$. In this setting, the test is not sound in the traditional sense; instead, the test is shown to satisfy a non-standard notion of soundness, which nevertheless is sufficient for the purposes of constructing a PCP. More specifically, in [25] it is proved that for all $\varepsilon > 0$ there is $r \in \mathbb{N}$ such that for sufficiently large ℓ and for tables T_1, T_2 as above, there are subspaces $Q \subseteq W \subseteq \mathbb{F}_q^n$ with $\dim(Q) + \operatorname{codim}(W) \leqslant r$ and a linear function $f: W \to \mathbb{F}_q$ such that

$$\Pr_{Q \subseteq L_1 \subseteq W}[T_1[L_1] \equiv f|_{L_1}] \geqslant \varepsilon'(\varepsilon) > 0.$$

We refer to the set

$$\{L \subseteq \mathbb{F}_q^n \mid \dim(L) = \ell_1, Q \subseteq L \subseteq W\}$$

as the zoom in of Q and zoom out of W. While this result is good for the purposes of 2-to-1 Games, the dependency between ℓ and ε (and thus, between the soundness and the alphabet size) is still not good enough for us.

Our low-degree test. It turns out that the proper setting of parameters for us is $\ell_2 = (1-\delta)\ell_1$ where $\delta > 0$ is a small constant. With these parameters, we are able to show that for $\varepsilon \geqslant q^{-(1-\delta')\ell_1}$ (where $\delta' = \delta'(\delta) > 0$ is a vanishing function of δ), if T_1, T_2 pass the test with probability at least ε, then there are subspaces $Q \subseteq W$ with $\dim(Q) + \operatorname{codim}(W) \leqslant r = r(\delta) \in \mathbb{N}$, and a linear function $f: W \to \mathbb{F}_q$ such that

$$\Pr_{Q \subseteq L_1 \subseteq W}[T_1[L_1] \equiv f|_{L_1}] \geqslant \varepsilon'(\varepsilon) = \Omega(\varepsilon).$$

This result is obtained from Lemma 2.2, which in turn relies on [16].

Working in the very small soundness regime of $\varepsilon \geqslant q^{-(1-\delta')\ell_1}$ entails with it many challenges, however. First, dealing with such small soundness requires us to use a strengthening of the global hypercontractivity result of [25] in the form of an optimal level d inequality due to Evra, Kindler and Lifshitz [16]. Second, in the context of [25], ε' could be any function of ε (and indeed it ends up being a polynomial function of ε). In the context of the current paper, it is crucial that $\varepsilon' = \varepsilon^{1+o(1)}$, as opposed to, say, $\varepsilon' = \varepsilon^{1.1}$. The reason is that, as we are dealing with very small ε, the result would be trivial for $\varepsilon' = \varepsilon^{1.1}$ and not useful towards the analysis of the PCP (as then ε' would be below the threshold $q^{-\ell_1}$ which represents the agreement a random linear function f has with T_1).

3.2 Getting List Decoding Bounds

As is usually the case in PCP reductions, we require a list decoding version for our low-degree test. Indeed, using a standard argument we are able to show that in the setting that $\ell_2 = (1-\delta)\ell_1$ and $\varepsilon \geqslant q^{(1-\delta')\ell_1}$, there is $r = r(\delta,\delta') \in \mathbb{N}$ such that for at least $q^{-\Theta(\ell_1)}$ fraction of subspaces $Q \subseteq \mathbb{F}_q^n$ of dimension r, there exists a subspace W with co-dimension at most r and $Q \subseteq W \subseteq \mathbb{F}_q^n$, as well as a linear function $f : W \to \mathbb{F}_q$, such that

$$\Pr_{Q \subseteq L_1 \subseteq W}[T_1[L_1] \equiv f|_{L_1}] \geqslant \varepsilon'(\varepsilon) = \Omega(\varepsilon). \tag{1}$$

This list decoding version theorem alone is not enough. In our PCP construction, we compose the inner PCP with an outer PCP (that we describe below), and analyzing the composition requires decoding global linear functions (from a list decoding version theorem as above) in a coordinated manner between two non communicating parties. Often times, the number of possible global functions that may be decoded is constant, in which case randomly sampling one among them often works. This is not the case for us, though: if (Q,W) and (Q',W') are distinct zoom-in and zoom-out pairs for which there are linear functions $f_{Q,W}$ and $f_{Q',W'}$ satisfying (1), then the functions $f_{Q,W}$ and $f_{Q',W'}$ could be completely different. Thus, to achieve a coordinated decoding procedure, we must:

(1) Facilitate a way for the two parties to agree on a zoom-in and zoom-out pair (Q,W) with noticeable probability.
(2) Show that for each (Q,W) there are at most $\mathrm{poly}(1/\varepsilon)$ functions $f_{Q,W}$ for which

$$\Pr_{Q \subseteq L_1 \subseteq W}[T_1[L_1] \equiv f_{Q,W}|_{L_1}] \geqslant \varepsilon'.$$

The second item is precisely the reason we need ε' to be $\varepsilon^{1+o(1)}$; any worse dependency, such as $\varepsilon' = \varepsilon^{1.1}$ would lead to the second item being false. We also remark that the number of functions being $\mathrm{poly}(1/\varepsilon)$ is important to us as well. There is some slack in this bound, but a weak quantitative bound such as $\exp(\exp(1/\varepsilon))$ would have been insufficient for some of our applications. Luckily, such bounds can be deduced from [19] for the case of linear functions.[4]

We now move onto the first item, in which we must facilitate a way for two non-communicating parties to agree on a zoom-in and zoom-out pair (Q,W). It turns out that agreeing on the zoom-in Q can be delegated to the outer PCP, and we can construct a sound outer PCP game in which the two parties are provided with a coordinated zoom-in Q. This works because in our list decoding theorem, the fraction of zoom-ins Q that work is significant. Coordinating zoom-outs is more difficult, and this is where much of the novelty in our analysis lies.

[4] In the case of higher degree functions (even quadratic functions) some bounds are known [6, 20] but they would not have been good enough for us.

3.3 Coordinating Zoom-outs

For the sake of simplicity and to focus on the main ideas, we ignore zoom-ins for now and assume that the list decoding statement holds with no Q. Thus, the list decoding theorem asserts that there exists a zoom-out W of constant co-dimension on which there is a global linear function. However, there could be many such zoom-outs W, say $W_1,\ldots,W_m$ and say all of them were of co-dimension $r = O_{\delta,\delta'}(1)$. If the number m were sufficiently large – say at least $q^{-\mathrm{poly}(\ell_1)}$ fraction of all co-dimension r subspaces – then we would have been able to coordinate them in the same way as we coordinate zoom-ins. If the number m were sufficiently small – say $m = q^{\mathrm{poly}(\ell_1)}$, then randomly guessing a zoom-out would work well enough. The main issue is that the number m is intermediate, say $m = q^{\sqrt{n}}$.

This issue had already appeared in [12, 24]. Therein, this issue is resolved by showing that if there are at least $m \geqslant q^{100\ell_1^2}$ zoom-outs $W_1,\ldots,W_m$ of co-dimension r, and linear functions $f_1,\ldots,f_m$ on $W_1,\ldots,W_m$ respectively such that

$$\Pr_{L \subseteq W_i}[T[L] \equiv f_i|_L] \geqslant \varepsilon'$$

for all i, then there exists a zoom out W of co-dimension *strictly less than* r and a linear function $f : W \to \mathbb{F}_q$ such that

$$\Pr_{L \subseteq W}[T[L] \equiv f|_L] \geqslant \Omega(\varepsilon'^{12}).$$

Thus, if there are too many zoom-outs of a certain co-dimension, then there is necessarily a zoom-out of smaller co-dimension that also works. In that case, the parties could go up to that co-dimension.

This result is not good enough for us, due to the polynomial gap between the agreement between and f_i's and F and the agreement between f an T. Indeed, in our range of parameters, ε'^{12} will be below the trivial threshold $q^{-\ell_1}$ which is the agreement a random linear function f has with T, and therefore the promise on the function f above is meaningless.

We resolve this issue by showing a stronger, essentially optimal version of the above assertion still holds. Formally, we prove:

Lemma 3.1. *For all $\delta > 0$, $r \in \mathbb{N}$ there is $C > 1$ such that the following holds for $\varepsilon' \geqslant q^{(1-\delta)\ell_1}$. Suppose that F is a table that assigns to each subspace L of dimension ℓ_1 a linear function, and suppose that there are at least $m \geqslant q^{C\ell_1}$ subspaces $W_1,\ldots,W_m$ of co-dimension r and linear functions $f_i : W_i \to \mathbb{F}_q$ such that*

$$\Pr_{L \subseteq W_i}[T[L] \equiv f_i|_L] \geqslant \varepsilon'$$

for all $i = 1,\ldots,m$. Then, there exists a zoom-out W of co-dimension strictly smaller than r and a linear function $f : W \to \mathbb{F}_q$ such that

$$\Pr_{L \subseteq W}[T[L] \equiv f|_L] \geqslant \Omega(\varepsilon').$$

We remark that our proof of Lemma 3.1 is very different from the arguments in [12] and is significantly more involved. Our proof

uses tools from [12, 24], tools from the analysis of the classical Plane vs Plane and direct product testing [23, 32, 39], global hypercontractivity [16] as well as Fourier analysis over the Grassmann graph.

3.4 The Outer PCP

Our outer PCP game is the outer PCP of [12, 24], which is a smooth parallel repetition of the equation versus variables game of Hastad [22] (or of [26] for the application to Quadratic Programming). As in there, we equip this game with the "advice" feature to facilitate zoom-in coordination (as discussed above). For the sake of completeness we elaborate on the construction of the outer PCP below.

We start with an instance of 3-Lin that has a large gap between the soundness and completeness. Namely, we start with an instance (X, E) of linear equations over $\mathbb{F}_q$ in which each equation has the form $ax_{i_1} + bx_{i_2} + cx_{i_3} = d$. It is known [22] that for all $\eta > 0$, it is NP-hard to distinguish between the following two cases:

(1) YES case: $\mathrm{val}(X, E) \geqslant 1 - \eta$.
(2) NO case: $\mathrm{val}(X, E) \leqslant \frac{1.1}{q}$.

Given the instance (X, E), we construct a 2-Prover-1-Round game, known as the smooth equation versus variable game with r-advice as follows. The verifier has a smoothness parameter $\beta > 0$ and picks a random equation e, say $ax_{i_1} + bx_{i_2} + cx_{i_3} = d$, from (X, E). Then:

(1) With probability $1-\beta$ the verifier takes $U = V = \{x_{i_1}, x_{i_2}, x_{i_3}\}$ and vectors $u_1 = v_1, \ldots, u_r = v_r \in \mathbb{F}_q^U$ sampled uniformly and independently.
(2) With probability β, the verifier sets $U = \{x_{i_1}, x_{i_2}, x_{i_3}\}$, chooses a set consisting of a single variable $V \subseteq U$ uniformly at random. The verifier picks $v_1, \ldots, v_r \in \mathbb{F}_q^V$ uniformly and independently and appends to each v_i the value 0 in the coordinates of $U \setminus V$ to get $u_1, \ldots, u_r$.

After that, the verifier sends U and $u_1, \ldots, u_r$ to the first prover and V and $v_1, \ldots, v_r$ to the second prover. The verifier expects to get from them $\mathbb{F}_q$ assignments to the variables in U and in V, and accepts if and only if these assignments are consistent, and furthermore the assignment to U satisfies the equation e.

Denoting the equation versus variable game by $^-$, it is easy to see that if $\mathrm{val}(X, E) \geqslant 1 - \eta$, then $\mathrm{val}(\Psi) \geqslant 1 - \eta$, and if $\mathrm{val}(X, E) \leqslant 1.1/q$, then $\mathrm{val}(\Psi) \leqslant 1 - \Omega(q^{-r}\beta)$. The gap between $1 - \eta$ and $1 - \Omega(q^{-r}\beta)$ is too weak for us, and thus we apply parallel repetition.

In the parallel repetition of the smooth equation versus variable game with advice, denoted by $\Psi^{\otimes k}$, the verifier picks k equations uniformly and independently $e_1, \ldots, e_k$, and picks U_i, $u_{1,i}, \ldots, u_{r,i}$ and V_i, $v_{1,i}, \ldots, v_{r,i}$ for each $i = 1, \ldots, k$ from e_i independently. Thus, the questions of the provers may be seen as $U = U_1 \cup \ldots \cup U_k$ and $V = V_1 \cup \ldots \cup V_k$ and their advice is $\vec{u}_j = (u_{j,1}, \ldots, u_{j,k}) \in \mathbb{F}_q^U$ for $j = 1, \ldots, r$ and $\vec{v}_j = (v_{j,1}, \ldots, v_{j,k}) \in \mathbb{F}_q^V$ for $j = 1, \ldots, r$ respectively. The verifier expects to get from the first prover a vector in $\mathbb{F}_q^U$ which specifies an $\mathbb{F}_q$ assignment to U, and from the second prover a vector in $\mathbb{F}_q^V$ specifying an $\mathbb{F}_q$ assignment to V. The verifier accepts if and only if these assignments are consistent and the assignment of the first prover satisfies all of $e_1, \ldots, e_k$. It is clear that if $\mathrm{val}(X, E) \geqslant 1 - \eta$, then $\mathrm{val}(\Psi^{\otimes m}) \geqslant 1 - k\eta$. Using the parallel repetition theorem of Rao [37] (albeit not in a completely trivial way) we argue that if $\mathrm{val}(X, E) \leqslant \frac{1.1}{q}$, then $\mathrm{val}(\Psi^{\otimes k}) \leqslant 2^{-\Omega(\beta q^{-r} k)}$. The game $\Psi^{\otimes k}$ is our outer PCP game.

Remark 3.2. *We remark that in the case of the Quadratic Programming application, we require a hardness result in which the completeness is very close to* 1 *in the form of Theorem 2.2. The differences between the reduction in that case and the reduction presented above are mostly minor, and amount to picking the parameters a bit differently. There is one significant difference in the analysis; we require a much sharper form of the "covering property" used in [12, 24], as elaborated on in Section 3.6*

3.5 Composing the Outer PCP and the Inner PCP Game

To compose the outer and inner PCPs, we take the outer PCP game, only keep the questions U to the first prover, and consider an induced 2-Prover-1-Round game on it. The alphabet is $\mathbb{F}_q^{3k}$: given a question U, the alphabet is the set of $\mathbb{F}_q$ assignment to the variables of U. There is a constraint between U and U' if there is a question V to the second prover such that $V \subseteq U \cap U'$. Denoting the assignments to U and U' by s_U and $s_{U'}$, the constraint between U and U' is that s_U satisfies all of the equations that form U, $s_{U'}$ satisfies all of the equations that form U', and s_U, $s_{U'}$ agree on $U \cap U'$.

The composition amounts to replacing each question U with a copy of our inner PCP. Namely, we identify between the question U and the space $\mathbb{F}_q^U$, and then replace U by a copy of the ℓ_2, ℓ_1 sub-spaces graph of $\mathbb{F}_q^U$. The answer s_U is naturally identified with the linear function $f_U(x) = \langle s_U, x \rangle$, which is then encoded by the sub-spaces encoding via tables of assignments $T_{1,U}$ and $T_{2,U}$.

The constraints of the composed PCP must check two things: (1) side conditions: the encoded vector s_U satisfies the equations of U, and (2) consistency: s_U and $s_{U'}$ agree on $U \cap U'$.

The first set of constraints is addressed by the folding technique, which we omit from this discussion. The second set of constraints is addressed by the ℓ_1 vs ℓ_2 subspace test, except that we have to modify it so that it works across blocks U and U'. This completes the description of the composition step of the outer PCP and the inner PCP, and thereby the description of our reduction.

Let us briefly describe the setting of parameters used to obtain Theorem 1.3. After fixing the δ, ε therein, we may take $q = 2$, choose δ' sufficiently small according to δ and ε, set $\ell_2 = (1 - \delta')\ell_1$, and finally take ℓ_1 sufficiently large. We must also choose k and β carefully to satisfy the covering property and completeness of the composed PCP, but omit further details from the current discussion. Altogether this yields alphabet size q^{ℓ_1} and soundness $q^{-(1-\varepsilon)\ell_1}$. We remark that the same tradeoff can be obtained with larger settings of q and this is indeed required for the application to hardness of approximating quadratic programming in Theorem 1.4.

3.6 The Covering Property

We end by briefly discussing the covering property. The covering property is an important feature of our outer PCP construction which enables the composition step to go through. The covering property first appeared in [27] and later more extensively in the

context of the 2-to-1 Games [12, 24]. To discuss the covering property, let $k \in \mathbb{N}$ be thought of as large, let $\beta \in (0,1)$ be thought of as $k^{-0.99}$ and consider sets $U_1, \ldots, U_k$ consisting of distinct element, each U_i has size 3 (in our context, U_i will be the set of variables in the ith equation the verifier chose). Let $U = U_1 \cup \ldots \cup U_k$, and consider the following two distributions over tuples in $\mathbb{F}_q^U$:

(1) Sample $x_1, \ldots, x_\ell \in \mathbb{F}_q^U$ uniformly.
(2) For each i independently, take $V_i = U_i$ with probability $1-\beta$ and otherwise take $V_i \subseteq U_i$ randomly of size 1, then set $V = V_1 \cup \ldots \cup V_k$. Sample $x_1, \ldots, x_\ell \in \mathbb{F}_q^V$ uniformly and lift them to points in $\mathbb{F}_q^U$ by appending 0's in $U \setminus V$. Output the lifted points.

In [24] it is shown that the two distributions above are $q^{3\ell}\beta\sqrt{k}$ close in statistical distance, which is good enough for the purposes of Theorem 1.3. However, this is not good enough for Theorem 1.4. [5] Carrying out a different analysis, we are able to show that the two distributions are close with better parameters and in a stronger sense: there exists a set E of ℓ tuples which has negligible measure in both distributions, such that each tuple not in E is assigned the same probability under the two distribution up to factor $(1+o(1))$. We are able to prove this statement provided that k is only slightly larger than $q^{2\ell}$.

The issue with the above two distributions is that they are actually far from each other if, say, $k = q^{1.9\ell}$. To see that, one can notice that the expected number of i's such that each one of $x_1, \ldots, x_\ell$ has the form $(a,0,0) \in \mathbb{F}_q^3$ on coordinates corresponding to U_i is very different. In the first distribution, this expectation is $\Theta(q^{-2\ell}k)$ which is less than 1, whereas in the second distribution it is at least $\beta k \geqslant k^{0.01}$.

To resolve this issue and to obtain nearly tight hardness in the Quadratic Programming application, we have to modify the distributions in the covering property so that (a) they will be close even if $k = q^{1.01\ell}$, and (b) we can still use these distributions in the composition step in our analysis of the PCP construction. Indeed, this is the route we take, and the two distributions we use are defined as follows:

(1) Sample $x_1, \ldots, x_\ell \in \mathbb{F}_q^U$ uniformly.
(2) For each i independently, take $V_i = U_i$ with probability $1-\beta$ and otherwise take $V_i \subseteq U_i$ randomly of size 1, then set $V = V_1 \cup \ldots \cup V_k$. Sample $x_1, \ldots, x_\ell \in \mathbb{F}_q^V$ uniformly, and let $w_i = 1_{U_i} \in \mathbb{F}_q^U$ be the vector that has 1 on coordinates of U_i and 0 everywhere else. Lift the points $x_1, \ldots, x_\ell$ to $x'_1, \ldots, x'_\ell \in \mathbb{F}_q^U$ by appending 0's in $U \setminus V$ and take $y_j = x_j + \sum_{i=1}^{k} \alpha_{i,j} w_i$ where $\alpha_{i,j}$ are independent random elements from $\mathbb{F}_q$. Output $y_1, \ldots, y_\ell$.

We show that for a suitable choice of k and β, these distributions are close even in the case that $k = q^{1.01\ell}$. [6] Indeed, as a sanity check one could count the expected number of appearances of blocks of the form $(0,a,0) \in \mathbb{F}_q^3$ and see they are very close ($q^{-2\ell}k$ versus $(1-\beta)q^{-2\ell}k + \beta k q^{-\ell}$). In this setting of parameters, k is roughly equal to the alphabet size – which can be made to be equal $(\log N)^{1-o(1)}$ under quasi-polynomial time reductions – it is sufficient to get the result of Theorem 1.4.

Remark 3.3. *We remark that a tight covering property is crucial for obtaining the tight hardness of approximation factor in Theorem 1.4. In the reduction from 2-Prover-1-Round games to Quadratic Programs, which is due to [1], the size of the resulting instance is exponential in the alphabet size and the soundness error remains roughly the same. In our case the alphabet size is roughly k hence the instance size is dominated by $N = 2^{\Theta(k^{1+o(1)})}$. If our analysis required $k = q^{C\ell}$, then even showing an optimal soundness of $q^{-(1-o(1))\ell}$ for the 2-Prover-1-Round game would only yield a factor of $(\log N)^{1/C-o(1)}$ hardness for quadratic programming.*

[5] The reason is that letting N be the size of the instance we produce, it holds that k is roughly logarithmic $\log N$ and q^ℓ is the alphabet size. To have small statistical distance, we must have $k \leqslant q^{6\ell}$, hence the soundness error could not go lower than $(\log N)^{-1/6}$.

[6] More speifically, one takes a small $c > 0$ and chooses $\beta = k^{2c/3-1}$, $k = q^{(1+c)\ell}$.

ACKNOWLEDGEMENTS

We thank the anonymous reviewers for their comments.

REFERENCES

[1] Sanjeev Arora, Eli Berger, Elad Hazan, Guy Kindler, and Muli Safra. 2005. On Non-Approximability for Quadratic Programs. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS 2005), 23-25 October 2005, Pittsburgh, PA, USA, Proceedings*. IEEE Computer Society, 206–215. https://doi.org/10.1109/SFCS.2005.57

[2] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. 1992. Proof Verification and Hardness of Approximation Problems. In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*. 14–23. https://doi.org/10.1109/SFCS.1992.267823

[3] Sanjeev Arora and Madhu Sudan. 2003. Improved Low-Degree Testing and its Applications. *Comb.* 23, 3 (2003), 365–426. https://doi.org/10.1007/S00493-003-0025-0

[4] Piotr Berman. 2000. A d/2 approximation for maximum weight independent set in d-claw free graphs. *Nordic Journal of Computing* 7, 3 (2000), 178–184.

[5] Amey Bhangale, Irit Dinur, and Inbal Livni Navon. 2017. Cube vs. Cube Low Degree Test. In *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA (LIPIcs, Vol. 67)*, Christos H. Papadimitriou (Ed.). Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 40:1–40:31. https://doi.org/10.4230/LIPICS.ITCS.2017.40

[6] Abhishek Bhowmick and Shachar Lovett. 2015. The List Decoding Radius of Reed-Muller Codes over Small Fields. In *Proceedings of the Forty-Seventh Annual ACM on Symposium on Theory of Computing, STOC 2015, Portland, OR, USA, June 14-17, 2015*, Rocco A. Servedio and Ronitt Rubinfeld (Eds.). ACM, 277–285. https://doi.org/10.1145/2746539.2746543

[7] Tanmoy Chakraborty, Julia Chuzhoy, and Sanjeev Khanna. 2008. Network design for vertex connectivity. In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, Cynthia Dwork (Ed.). ACM, 167–176. https://doi.org/10.1145/1374376.1374403

[8] Siu On Chan. 2016. Approximation resistance from pairwise-independent subgroups. *Journal of the ACM (JACM)* 63, 3 (2016), 1–32.

[9] Moses Charikar and Anthony Wirth. 2004. Maximizing Quadratic Programs: Extending Grothendieck's Inequality. In *45th Symposium on Foundations of Computer Science (FOCS 2004), 17-19 October 2004, Rome, Italy, Proceedings*. IEEE Computer Society, 54–60. https://doi.org/10.1109/FOCS.2004.39

[10] Julia Chuzhoy, Mina Dalirrooyfard, Vadim Grinberg, and Zihan Tan. 2022. A New Conjecture on Hardness of Low-Degree 2-CSP's with Implications to Hardness of Densest k-Subgraph and Other Problems. *CoRR* abs/2211.05906 (2022). https://doi.org/10.48550/ARXIV.2211.05906 arXiv:2211.05906

[11] Marek Cygan, Fabrizio Grandoni, and Monaldo Mastrolilli. 2013. How to Sell Hyperedges: The Hypermatching Assignment Problem. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, Sanjeev Khanna (Ed.). SIAM, 342–351. https://doi.org/10.1137/1.9781611973105.25

[12] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. 2018. Towards a proof of the 2-to-1 games conjecture?. In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, Ilias Diakonikolas, David Kempe, and Monika Henzinger (Eds.). ACM, 376–389. https://doi.org/10.1145/3188745.3188804

[13] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. 2021. On non-optimally expanding sets in Grassmann graphs. *Israel Journal of Mathematics* 243, 1 (2021), 377–420.

[14] Pavel Dvořák, Andreas Emil Feldmann, Ashutosh Rai, and Paweł Rzążewski. 2023. Parameterized inapproximability of independent set in H-free graphs. *Algorithmica* 85, 4 (2023), 902–928.

[15] David Ellis, Guy Kindler, and Noam Lifshitz. 2023. An Analogue of Bonami's Lemma for Functions on Spaces of Linear Maps, and 2-2 Games. In *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, Barna Saha and Rocco A. Servedio (Eds.). ACM, 656–660. https://doi.org/10.1145/3564246.3585116

[16] Shai Evra, Guy Kindler, and Noam Lifshitz. 2024+. HYPERCONTRACTIVITY FOR GLOBAL FUNCTIONS ON THE GENERAL LINEAR GROUP OVER A FINITE FIELD. (2024+).

[17] Uriel Feige. 1998. A Threshold of ln *n* for Approximating Set Cover. *J. ACM* 45, 4 (1998), 634–652. https://doi.org/10.1145/285055.285059

[18] Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. 1991. Approximating Clique is Almost NP-Complete (Preliminary Version). In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*. 2–12. https://doi.org/10.1109/SFCS.1991.185341

[19] Oded Goldreich, Ronitt Rubinfeld, and Madhu Sudan. 2000. Learning Polynomials with Queries: The Highly Noisy Case. *SIAM J. Discret. Math.* 13, 4 (2000), 535–570. https://doi.org/10.1137/S0895480198344540

[20] Parikshit Gopalan. 2013. A Fourier-Analytic Approach to Reed-Muller Decoding. *IEEE Trans. Inf. Theory* 59, 11 (2013), 7747–7760. https://doi.org/10.1109/TIT.2013.2274007

[21] Johan Håstad. 1996. Clique is Hard to Approximate Within $n^{1\text{-epsilon}}$. In *37th Annual Symposium on Foundations of Computer Science, FOCS '96, Burlington, Vermont, USA, 14-16 October, 1996*. IEEE Computer Society, 627–636. https://doi.org/10.1109/SFCS.1996.548522

[22] Johan Håstad. 2001. Some optimal inapproximability results. *J. ACM* 48, 4 (2001), 798–859. https://doi.org/10.1145/502090.502098

[23] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. 2012. New Direct-Product Testers and 2-Query PCPs. *SIAM J. Comput.* 41, 6 (2012), 1722–1768. https://doi.org/10.1137/09077299X

[24] Subhash Khot, Dor Minzer, and Muli Safra. 2017. On independent sets, 2-to-2 games, and Grassmann graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, Hamed Hatami, Pierre McKenzie, and Valerie King (Eds.). ACM, 576–589. https://doi.org/10.1145/3055399.3055432

[25] Subhash Khot, Dor Minzer, and Muli Safra. 2023. Pseudorandom sets in Grassmann graph have near-perfect expansion. *Annals of Mathematics* 198, 1 (2023), 1–92.

[26] Subhash Khot and Ashok Kumar Ponnuswami. 2006. Better Inapproximability Results for MaxClique, Chromatic Number and Min-3Lin-Deletion. In *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I (Lecture Notes in Computer Science, Vol. 4051)*, Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener (Eds.). Springer, 226–237. https://doi.org/10.1007/11786986_21

[27] Subhash Khot and Muli Safra. 2013. A Two-Prover One-Round Game with Strong Soundness. *Theory Comput.* 9 (2013), 863–887. https://doi.org/10.4086/toc.2013.v009a028

[28] Bundit Laekhanukit. 2014. Parameters of Two-Prover-One-Round Game and The Hardness of Connectivity Problems. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, Portland, Oregon, USA, January 5-7, 2014*, Chandra Chekuri (Ed.). SIAM, 1626–1643. https://doi.org/10.1137/1.9781611973402.118

[29] Euiwoong Lee and Pasin Manurangsi. 2023. Hardness of Approximating Bounded-Degree Max 2-CSP and Independent Set on k-Claw-Free Graphs. *CoRR* abs/2309.04099 (2023). https://doi.org/10.48550/arXiv.2309.04099 arXiv:2309.04099

[30] Pasin Manurangsi. 2019. A note on degree vs gap of Min-Rep Label Cover and improved inapproximability for connectivity problems. *Inf. Process. Lett.* 145 (2019), 24–29. https://doi.org/10.1016/J.IPL.2018.08.007

[31] Alexandre Megretski. 2001. Relaxations of quadratic programs in operator theory and system analysis. In *Systems, Approximation, Singular Integral Operators, and Related Topics: International Workshop on Operator Theory and Applications, IWOTA 2000*. Springer, 365–392.

[32] Dor Minzer and Kai Zheng. 2023. Approaching the Soundness Barrier: A Near Optimal Analysis of the Cube versus Cube Test. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, Nikhil Bansal and Viswanath Nagarajan (Eds.). SIAM, 2761–2776. https://doi.org/10.1137/1.9781611977554.ch104

[33] Dor Minzer and Kai Zhe Zheng. 2024. Near Optimal Alphabet-Soundness Tradeoff PCPs. *Electron. Colloquium Comput. Complex.* (2024), TR24–027. https://eccc.weizmann.ac.il/report/2024/027

[34] Dana Moshkovitz and Ran Raz. 2010. Two-query PCP with subconstant error. *J. ACM* 57, 5 (2010), 29:1–29:29. https://doi.org/10.1145/1754399.1754402

[35] Arkadi Nemirovski, Cornelis Roos, and Tamás Terlaky. 1999. On maximization of quadratic form over intersection of ellipsoids with common center. *Mathematical programming* 86, 3 (1999), 463–473.

[36] Zeev Nutov. 2012. Approximating minimum-cost connectivity problems via uncrossable bifamilies. *ACM Transactions on Algorithms (TALG)* 9, 1 (2012), 1–16.

[37] Anup Rao. 2011. Parallel Repetition in Projection Games and a Concentration Bound. *SIAM J. Comput.* 40, 6 (2011), 1871–1891. https://doi.org/10.1137/080734042

[38] Ran Raz. 1998. A Parallel Repetition Theorem. *SIAM J. Comput.* 27, 3 (1998), 763–803. https://doi.org/10.1137/S0097539795280895

[39] Ran Raz and Shmuel Safra. 1997. A sub-constant error-probability low-degree test, and a subconstant error-probability pcp characterization of np. In *Proceedings of the twenty-ninth annual ACM Symposium on Theory of Computing*. 475–484.

[40] Theophile Thiery and Justin Ward. 2023. An Improved Approximation for Maximum Weighted *k*-Set Packing. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, Nikhil Bansal and Viswanath Nagarajan (Eds.). SIAM, 1138–1162. https://doi.org/10.1137/1.9781611977554.CH42

[41] Luca Trevisan. 2014. Inapproximability of combinatorial optimization problems. *Paradigms of Combinatorial Optimization: Problems and New Approaches* (2014), 381–434.

Received 11-NOV-2023; accepted 2024-02-11