Electronic Colloquium on Computational Complexity, Report No. 27 (2024)

Near Optimal Alphabet-Soundness Tradeoff PCPs

Dor Minzer* Kai Zhe Zheng†

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 [Cha16]. 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 and Safra [KS13] and nearly matches the performance of the best known approximation algorithm [Meg01, NRT99, CW04] 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 and Manurangsi [LM23].
3. Improved hardness results for connectivity problems: using results of Laekhanukit [Lae14] and Manurangsi [Man19], 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.

1 Introduction

The PCP theorem is a fundamental result in theoretical computer science with many equivalent formulations [FGL$^+$91, AS97, ALM$^+$92]. 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 [Hås01, Has99, Fei98, KP06] (see [Tre14] for a survey). To discuss these parameters, it is often convenient to view the PCP from via 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 $\phi_e : \Sigma_L \to \Sigma_R$.*

*Department of Mathematics, Massachusetts Institute of Technology, Cambridge, USA. Supported by NSF CCF award 2227876 and NSF CAREER award 2239160.

†Department of Mathematics, Massachusetts Institute of Technology, Cambridge, USA. Supported by the NSF GRFP DGE-2141064.

[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.

ISSN 1433-8092

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,*

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

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 agree beforehand on a strategy; after that period they are not allowed to communicate. The verifier then picks a random edge from the 2-Prover-1-Round game $e = (u, v)$, sends u to the first prover, v to the second prover, gets a label from each one of them and checks that the labels satisfy the constraint ϕ_e. If so, 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 provers' strategies. This view will be useful for us later.

In the language of 2-Prover-1-Round Games, the majority of hardness of approximation results are proved by combining the basic PCP theorem [FGL+91, AS97, ALM+92] with Raz's parallel repetition theorem [Raz98], which together imply the following result:

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:* $\mathsf{val}(\Psi) = 1$.

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

For many applications, one only requires that the soundness of the PCP is small. Namely, that $\mathsf{val}(\Psi)$ is arbitrarily small in the "NO case". For certain applications however, more is required: not only must the soundness be small – but it must also be small in terms of the alphabet size. The tradeoff between the soundness 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 [KS13] showed that Theorem 1.2 holds for $\gamma = 1/6$ with imperfect completeness (i.e., $\mathsf{val}(\Psi) \geqslant 1 - o(1)$ instead of $\mathsf{val}(\Psi) = 1$ in the YES case). The result of Khot and Safra was improved by Chan [Cha16], who showed (using a completely different set of techniques) that Theorem 1.2 holds for $\gamma = 1/2 - o(1)$, again with imperfect completeness.

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

Theorem 1.3 shows a near optimal tradeoff between the alphabet of a PCP and the alphabet of a PCP, improving upon the result of Chan [Cha16]. 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 alphabet size, which our construction achieves. Other applications require the underlying constraint graph to be bounded-degree bi-regular graph, which our construction also achieves (after mild modifications; see Theorem 7.1).

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 $Q(x) = \sum_{i,j=1}^{n} a_{i,j}x_ix_j$ 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 [Meg01, NRT99, CW04], and is known to be quasi-NP-hard to approximate within factor $(\log n)^{1/6-o(1)}$ [ABH+05, KS13]. 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 [ABH+05]. This connections requires a good tradeoff between the alphabet size, the soundness of the PCP, and the size of the PCP. Fortunately, the construction in Theorem 1.4 has a sufficiently good tradeoff between the alphabet size. [2]

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.

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

[2] We remark that the result of Chan [Cha16] does not achieve a good enough tradeoff 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.

assuming the Unique-Games Conjecture [LM23]. 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 [LM23], 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 $\frac{k}{2}$ [Ber00, TW23], and a quasi-polynomial time approximation algorithm within factor $\left(\frac{1}{3} + o(1)\right) k$ [CGM13]. As in [LM23], using ideas from [DFRR23] 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 [LM23] 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 [Lae14] and the improvements by Manurangsi [Man19], Theorem 1.3 implies improved hardness of approximation results for several graph connectivitiy problems. More specifically, Theorem 1.3 combined with the results of [Man19] implies improvements to each one of the results outlined in table 1 in [Lae14] 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 [Man19]. We briefly discuss the Rooted k-Connectivity Problem, but defer the reader to [Lae14] 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 [Nut12].

Using the ideas of [Lae14], 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 [CCK08], 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.

1.2 Our Techniques

Theorem 1.3 is proved via a composition of 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 [KS13, KMS17] while the inner PCP is based on the Grassmann graph [KMS17, DKK$^+$18, DKK$^+$21, KMS23].

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. 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 [KMS17, DKK$^+$18, DKK$^+$21, KMS23], 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 [EKL23b]. In addition, our analysis incorporates ideas from the plane versus plane test and direct product testing [RS97, IKW12, MZ23], from classical PCP theory [KS13], as well as from error correcting codes [GRS00]. All of these tools are necessary to prove our main technical statement – Lemma 1.7 below – which is a combinatorial statement that may be of independent interest.

We now elaborate on each one of the components separately.

1.2.1 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 $f\colon \mathbb{F}_q^n \to \mathbb{F}_q$ 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 [RS97]. 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 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 [MR10, BDN17, MZ23]. Currently, the best known analysis for the Plane vs Plane test [MR10] shows that one may take $c = 1/8$. Better analysis is known for higher dimensional encoding [BDN17, MZ23], and for the 3-dimensional version of it a near optimal soundness result is known [MZ23].

Low degree tests in this paper. In the context of the current paper, we wish to encode linear functions $f\colon \mathbb{F}_q^n \to \mathbb{F}_q$, 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 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 $f_i \colon W_i \to \mathbb{F}_q$ 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 setting of the parameters would give alphabet size $R = q^{\ell_1}$ and (possibly) soundness $\Theta(1/\sqrt{R})$. 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 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 [KMS23] 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 $\mathsf{dim}(Q) + \mathsf{codim}(W) \leqslant r$ and a linear function $f \colon W \to \mathbb{F}_q$ 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

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

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 $\mathsf{dim}(Q) + \mathsf{codim}(W) \leqslant r = r(\delta) \in \mathbb{N}$, and a linear function $f\colon W \to \mathbb{F}_q$ such that

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

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 [KMS23] in the form of an optimal level d inequality due to Evra, Kindler and Lifshitz [EKL23b]. Second, in the context of [KMS23], 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).

1.2.2 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 $f\colon W \to \mathbb{F}_q$, 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 $\mathsf{poly}(1/\varepsilon)$ 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 $\mathsf{poly}(1/\varepsilon)$

is important to us as well. There is some slack in this bound, but a weak quantitative bound such as $\mathsf{exp}(\mathsf{exp}(1/\varepsilon))$ would have been insufficient for some of our applications. Luckily, such bounds can be deduced from [GRS00] for the case of linear functions.[3]

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.

1.2.3 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 $q^{-\mathsf{poly}(\ell_1)}$ 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 $m = q^{\mathsf{poly}(\ell_1)}$, 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 [KMS17, DKK$^+$18]. 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 $f\colon W \to \mathbb{F}_q$ 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 1.7. *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 $f_i\colon W_i \to \mathbb{F}_q$ such that*

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

[3] In the case of higher degree functions (even quadratic functions) some bounds are known [Gop10, BL15] but they would not have been good enough for us.

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 defer a detailed discussion about Lemma 1.7 and its proof to Section 8, but remark that our proof of Lemma 1.7 is very different from the arguments in [DKK⁺18] and is significantly more involved. Our proof uses tools from [KMS17, DKK⁺18], tools from the analysis of the classical Plane vs Plane and direct product testing [RS97, IKW12, MZ23], global hypercontractivity [EKL23b] as well as Fourier analysis over the Grassmann graph.

1.2.4 The Outer PCP

Our outer PCP game is the outer PCP of [KMS17, DKK⁺18], which is a smooth parallel repetition of the equation versus variables game of Hastad [Hås01] (or of [KP06] 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 [Hås01] that for all $\eta > 0$, it is NP-hard to distinguish between the following two cases:

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

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 Ψ, it is easy to see that if $\mathsf{val}(X, E) \geqslant 1 - \eta$, then $\mathsf{val}(\Psi) \geqslant 1 - \eta$, and if $\mathsf{val}(X, E) \leqslant 1.1/q$, then $\mathsf{val}(\Psi) \leqslant 1 - \Omega(q^{-r}\beta)$. 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 $\mathsf{val}(X, E) \geqslant 1 - \eta$, then $\mathsf{val}(\Psi^{\otimes m}) \geqslant 1 - k\eta$. Using the parallel repetition theorem of Rao [Rao08] (albeit not in a completely trivial way) we argue that if $\mathsf{val}(X, E) \leqslant \frac{1.1}{q}$, then $\mathsf{val}(\Psi^{\otimes k}) \leqslant 2^{-\Omega(\beta q^{-r} k)}$. The game $\Psi^{\otimes k}$ is our outer PCP game.

Remark 1.8. *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 2.1. 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 [KMS17, DKK*$^{+}$*18], as elaborated on in Section 1.2.6*

1.2.5 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 $\mathbb{F}_q^{3k}$, that given a question U specifies an $\mathbb{F}_q$ 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 that: (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 other PCP and the inner PCP, and thereby the description of our reduction.

1.2.6 The Covering Property

We end this introductory section by 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 [KS13] and later more extensively in the context of the 2-to-1 Games [KMS17, DKK^{+}18]. 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 [KMS17] 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.

[4] 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 go all the way through 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 a suitable chose of k and β gives that these distributions are close even in the case that $k = q^{1.01\ell}$. [5] Indeed, as a sanity check one could count the expected number of appearances of blocks of the form $(0, a, 0) \in \mathbb{F}^3$ 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 1.9. *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 [ABH+05], the size of the resulting instance is exponential in the alphabet size and the soundness 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.*

2 Preliminaries

2.1 The Grassmann Graph

In this section we present the Grassmann graph and some Fourier analytic tools on it that are required for our analysis of the inner PCP. Throughout this section, we fix parameters n, ℓ with $1 \ll \ell \ll n$, and a prime power q.

[4] 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 could not go lower than $(\log N)^{-1/6}$.

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

2.1.1 Basic Definitions

The Grassmann graph $\mathsf{Grass}_q(n, \ell)$ 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 $\dim(L \cap L') = \ell - 1$.

At times we will have a vector space V over $\mathbb{F}_q$, and thus we may identify V with $\mathbb{F}_q^n$ and work with the Grassmann graph on the ℓ-dimensional subspaces $L \subseteq V$. We may also use $\mathsf{Grass}_q(V, \ell)$ to denote this graph, which is isomporphic to $\mathsf{Grass}_q(n, \ell)$ if $\dim(V) = n$. Abusing notation, we also use $\mathsf{Grass}_q(n, \ell)$ to denote the set of all ℓ-dimensional subspaces in $\mathbb{F}_q^n$. Throughout, we denote by $L_2(\mathsf{Grass}_q(n, \ell))$ the set of complex valued functions $F\colon \begin{bmatrix} \mathbb{F}_q^n \\ \ell \end{bmatrix}_q \to \mathbb{C}$.

The number of ℓ-dimensional subspaces of $\mathbb{F}_q^n$ is counted by the Gaussian binomial coefficient, $\begin{bmatrix} n \\ \ell \end{bmatrix}_q$. The following standard fact gives a formula for the Gaussian binomial coefficients, and we omit the proof.

Fact 2.1. *Suppose $1 \leqslant \ell \leqslant \frac{n}{2}$, then the number of vertices in $\mathsf{Grass}_q(n, \ell)$ is given by*

$$\begin{bmatrix} n \\ \ell \end{bmatrix}_q = \prod_{i=0}^{\ell-1} \frac{q^n - q^i}{q^\ell - q^i}.$$

Abusing notations, we denote by $\begin{bmatrix} V \\ \ell \end{bmatrix}_q$ the set of ℓ dimensional subspaces of V.

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

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

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

2.1.2 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. [6] Let $\mu(F) = \mathbb{E}_{L \in \mathsf{Grass}_q(n,\ell)}[F(L)]$ (for indicator functions, this is simply the measure of the indicated set). For subspaces $Q \subseteq W \subseteq \mathbb{F}_q^n$, define

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

[6] 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.

Definition 2.2. *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 $\dim(Q) + \operatorname{codim}(W) = r$, we have*

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

We will often say that a set $S \subseteq \mathsf{Grass}_q(n,\ell)$ 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 [EKL23a, EKL23b].

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

Lemma 2.3. *Let $F\colon \mathsf{Grass}_q(n,2\ell) \to \{0,1\}$ and $G\colon \mathsf{Grass}_q(n,2(1-\delta)\ell) \to \{0,1\}$ be Boolean functions such that $\mathbb{E}_L[F(L)] = \alpha$, $\mathbb{E}_R[G(R)] = \beta$, and suppose that F is (r,ε) pseudo-random. Then for all $t \geqslant 4$ that are powers of 2,*

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

Proof. Deferred to Section A ☐

We also need the following lemma, asserting that if a not-too-small set S is highly pseudo-random, then its density remains nearly the same on all zoom-ins.

Lemma 2.4. *For all $\xi > 0$, the following holds for sufficiently large ℓ. Suppose that $\ell' \geqslant \frac{\xi}{3}\ell$, $\delta_2 = \frac{\xi}{100}$, and let $V^\star$ be a subspace such that $\dim(V^\star) \geqslant \ell'^2$. Let $\mathcal{L}^\star \subseteq \mathsf{Grass}_q(V^\star,\ell')$ have measure $\mu(\mathcal{L}^\star) = \eta \geqslant q^{-2\ell}$ and set $Z = \{z \in V^\star \mid |\mu_z(\mathcal{L}^\star) - \eta| \leqslant \frac{\eta}{10}\}$. If $\mathcal{L}^\star$ is $(1, q^{\delta_2\ell/100})$-pseudo-random, then*

$$|Z| \geqslant \left(1 - q^{\frac{\ell'}{2}}\right)|V^\star|.$$

Proof. The proof is deferred to Appendix A.4 ☐

At times we will also use the term *global* to refer to sets whose indicator functions are pseudo-random. That is, we say that a set is (r,ε)-global if its indicator is (r,ε)-pseudo-random.

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 $\mathsf{3Lin}$ and the $\mathsf{Gap3Lin}$ problem.

Definition 2.5. *For a prime power q, an instance of $\mathsf{3Lin}$ 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 $\mathsf{Gap3Lin}[c, s]$ 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 $\mathsf{Gap3Lin}[c, s]$ 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 [Hås01]. 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 $\mathsf{Gap3Lin}[1-\varepsilon, 1/q+\varepsilon]$ 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,* $\mathsf{Gap3Lin}\,[1-\eta, s]$ *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 [KP06]. 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 a* $\mathsf{Gap3Lin}[1-\eta, 1-\varepsilon]$ *instance with 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 Outer PCP

In this section, we describe our outer PCP game. In short, our outer PCP is a smooth parallel repetition of the variable versus equation game with advice. This outer PCP was first considered in [KS13] without the advice feature, and then in [KMS17] with the advice feature.

3.1 The Outer PCP construction

Let $\varepsilon_1 < \varepsilon_2$ be parameters that determine the completeness and the soundness our Gap3Lin. Our reduction starts with the $\mathsf{Gap3Lin}[1-\varepsilon_1, 1-\varepsilon_2]$ problem, and we fix an instance of it (X, E) for the rest of this section. Our presentation is gradual, and we begin by presenting the basic Variable versus Equation Game. We then equip it with the additional features of smoothness and advice.

3.1.1 The Variable versus Equation Game

We first convert the instance (X, E) into an instance of 2-Prover-1-Round Games, and it will be convenient for us to describe it in the active view with a verifier and 2 provers.

In the Variable versus Equation game, the verifier picks an equation $e \in E$ uniformly at random, and then chooses a random variable $x \in e$. The verifier sends the question e, i.e. the three variables appearing in e, to the first prover, and sends the variable x to the second prover. The provers are expected to answer with assignments to their received variables, and the verifier accepts if and only if the two assignments agree on x and the first prover's assignment satisfies the equation e. If the verifier accepts then we also say that the provers pass. This game has the following completeness and soundness, which are both easy to see (we omit the formal proof):

1. **Completeness:** If (X, E) has an assignment satisfying $1 - \varepsilon$-fraction of the equations, then the prover's have a strategy that passes with probability at least $1 - \varepsilon$.
2. **Soundness:** If (X, Eq) has no assignment satisfying more than $1 - \varepsilon$-fraction of the equations, then the prover's can pass with probability at most $1 - \frac{\varepsilon}{3}$.

3.1.2 The Smooth Equation versus Variable Game

We next describe a smooth version of the Variable versus Equation game. In this game, the verifier has a parameter $\beta \in (0, 1]$, and it proceeds as follows:

1. The verifier chooses an equation $e \in E$ uniformly, and lets U be the set of variables in e.
2. With probability $1 - \beta$, the prover chooses $V = U$. With probability β, the prover chooses $V \subseteq U$ randomly of size 1.
3. The verifier sends U to the first prover, and V to the second prover.
4. The provers respond with assignments to the variables they receive, and the verifier accepts if and only if their assignments agree on V and the assignment to U satisfies the equation e.

The smooth Variable versus Equation game has the following completeness and soundness property, which are again easily seen to hold (we omit the formal proof).

1. **Completeness:** If (X, E) has an assignment satisfying $1-\varepsilon$ fraction of the equations, then the provers have a strategy that passes with probability at least $1 - \varepsilon$.
2. **Soundness:** If (X, Eq) has no assignment satisfying more than $1 - \varepsilon$ fraction of the equations, then the provers can pass with probability at most $1 - \frac{\beta\varepsilon}{3}$.

3.1.3 The Smooth Equation versus Variable Game with Advice

Next, we introduce the feature of advice into the smooth Variable versus Equation Game. This "advice" acts as shared randomness which may help the provers in their strategy; we show though that it does not considerably change the soundness. The game is denoted by $G_{\beta,r}$ for $\beta \in (0, 1]$ and $r \in \mathbb{N}$, and proceeds as follows:

1. The verifier chooses an equation $e \in E$ uniformly, and lets U be the set of variables in e.
2. With probability $1 - \beta$, the verifier chooses $V = U$. With probability β, the verifier chooses $V \subseteq U$ randomly of size 1.

3. The verifier picks vectors $v_1, \dots, v_r \in \mathbb{F}_q^V$ uniformly and independently. If $U = V$ the verifier takes $u_i = v_i$ for all i, and otherwise the verifier takes the vectors $u_1, \dots, u_r \in \mathbb{F}_q^U$ where for all $i = 1, \dots, r$, the vector u_i agrees with v_i on the coordinate of V, and is 0 in the coordinates of $U \setminus V$.

4. The verifier sends U and $u_1, \dots, u_r$ to the first prover, and V and $v_1, \dots, v_r$ to the second prover.

5. The provers respond with assignments to the variables they receive, and the verifier accepts if and only if their assignments agree on V and the assignment to U satisfies the equation e.

Below we state the completeness and soundness of this game:

1. **Completeness:** If (X, E) has an assignment satisfying $1-\varepsilon$ fraction of the equations, then the provers have a strategy that passes with probability at least $1 - \varepsilon$. This is easy to see.

2. **Soundness:** If (X, Eq) has no assignment satisfying more than $1 - \varepsilon$ fraction of the equations, then the provers can pass with probability at most $1 - \frac{q^{-r}\beta\varepsilon}{3}$. Indeed, suppose that the provers can win the game with probability at least $1 - \eta$. Note that with probability at least βq^{-r} it holds that $U \neq V$ and all the vectors $u_1, \dots, u_r$ and $v_1, \dots, v_r$ are all 0, in which case the provers play the standard equation versus variable game. Thus, the provers' strategy wins in the latter game with probability at least $1 - \frac{\eta}{q^{-r}\beta} > 1 - \frac{\varepsilon}{3}$, and contradiction.

3.1.4 Parallel Repetition of the Smooth Equation versus Variable Game with Advice

Finally, our Outer PCP is then the k-fold parallel repetition of $G_{\beta,r}$, which we denote by $G_{\beta,r}^{\otimes k}$. Below is a full description of it:

1. The verifier chooses equations $e_1, \dots, e_k \in E$ uniformly and independently, and lets U_i be the set of variables in e_i.

2. For each i independently, with probability $1 - \beta$, the verifier chooses $V_i = U_i$. With probability β, the verifier chooses $V_i \subseteq U_i$ randomly of size 1.

3. For each i independently, the verifier picks a vectors $v_1^i, \dots, v_r^i \in \mathbb{F}_q^V$ uniformly and independently. If $U_i = V_i$ the verifier takes $u_j^i = v_j^i$ for $j = 1, \dots, r$, and otherwise the verifier takes the vectors $u_1^i, \dots, u_r^i \in \mathbb{F}_q^U$ where for all $j = 1, \dots, r$, the vector u_j^i agrees with v_j^i on the coordinate of V_i, and is 0 in the coordinates of $U_i \setminus V_i$.

4. The verifier sets $U = \bigcup_{i=1}^k U_i$ and $u_j = (u_j^1, \dots, u_j^k)$ for each $j = 1, \dots r$, and $V = \cup_{i=1}^k V_i$ and $v_j = (v_j^1, \dots, v_j^k)$ for each $j = 1, \dots, r$. The verifier sends U and $u_1, \dots, u_r$ to the first prover, and V and $v_1, \dots, v_r$ to the second prover.

5. The provers respond with assignments to the variables they receive, and the verifier accepts if and only if their assignments agree on V and the assignment to U satisfies the equations $e_1, \dots, e_k$.

Next, we state the completeness and the soundness of the game $G_{\beta,r}^{\otimes k}$, and we begin with its completeness.

Claim 3.1. *If (X, E) has an assignment satisfying at least $1 - \varepsilon$ of the equations, then the provers can win $G_{\beta,r}^{\otimes k}$ with probability at least $1 - k\varepsilon$.*

Proof. Let A be an assignment that satisfies at least $1-\varepsilon$ fraction of the equations in E, and consider the strategy of the provers that assigns their variables according to A. Note that whenver each one of the equations $e_1,\ldots,e_k$ the verifier chose is satisfied by A, the verifier accepts. By the union bound, the probability this happens is at least $1-k\varepsilon$. $\square$

Next, we establish the soundness of the game $G_{\beta,r}^{\otimes k}$.

Claim 3.2. *If there is no assignment to (X,E) satisfying at least $1-\varepsilon$ of the equations, then the provers can win $G_{\beta,r}^{\otimes k}$ with probability at most $2^{-\Omega(\varepsilon^2 q^{-r}\beta k)}$.*

Proof. We appeal to the parallel repetition theorem for projection games of Rao [Rao08], but we have to do so carefully. That theorem states that if Ψ is a 2-Prover-1-Round game with $\mathsf{val}(\Psi)\leqslant 1-\eta$, then $\mathsf{val}(\Psi^{\otimes k})\leqslant 2^{-\Omega(\eta^2 k)}$. We cannot apply the theorem directly on $G_{\beta,r}$ (as the square is too costly for us). Instead, we consider the game $\Psi=G_{\beta,r}^{\frac{q^r}{\beta}}$ and note that it has value bounded away from 1.

Write $\mathsf{val}(\Psi)=1-\eta$. Note that probability at least 0.1 there exists at least a single coordinate i in which $U_i\neq V_i$ and all of the advice vectors $v_1^i,\ldots,v_r^i$ and $u_1^i,\ldots,u_r^i$ are all 0. Thus, there exists a coordinate i and a fixing for the questions of the provers outside i so that the answers of the players to the ith coordinate win the standard equation versus variable game with probability at least $1-10\eta$. It follows that $1-10\eta\leqslant 1-\frac{\varepsilon}{3}$, and so $\eta\geqslant\frac{\varepsilon}{30}$.

We conclude from Rao's parallel repetition theorem that

$$\mathsf{val}(G_{\beta,r}^{\otimes k})=\mathsf{val}(\Psi^{\otimes\frac{k}{q^r/\beta}})\leqslant 2^{-\Omega(\varepsilon^2 q^{-r}\beta k)}. \qquad \square$$

Viewing the advice as subspaces. Due to the fact that each variable appears in at most $O(1)$ equations, it can easily be seen that with probability $1-O(k^2/n)$, all variables in $e_1,\ldots,e_k$ are distinct. In that case, note that the r vectors of advice to the second prover, $v_1,\ldots,v_r\in\mathbb{F}_q^V$, are uniform, and the second prover may consider their span Q_V. Note that the distribution of Q_V is that of a uniform r dimensional subspace of $\mathbb{F}_q^V$. As for the second prover, the vectors $u_1,\ldots,u_r\in\mathbb{F}_q^U$ are not uniformly distributed. Nevertheless, as shown by the covering property from [KS13, KMS17] (and presented below), the distribution of $u_1,\ldots,u_r$ is close to uniform over r-tuple of vectors from $\mathbb{F}_q^U$. Thus, the first prover can also take their span, call it Q_U, and think of it as a random r-dimensional subspace of $\mathbb{F}_q^U$ (which is highly correlated to Q_V).

4 The Composed PCP Construction

In this section we describe the final PCP construction, which is a composition of the outer PCP from Section 3 with the inner PCP based on the Grassmann consistency test.

4.1 The Underlying Graph

Our instance of 2-Prover-1-Round Games starts from an instance (X,Eq) of $\mathsf{Gap3Lin}$. Consider the game $G_{\beta,r}^{\otimes k}$ from Section 3, and let $\mathcal{U}$ denote the set of questions asked to the first prover. Thus $\mathcal{U}$ consists of all k-tuples of equations $U=(e_1,\ldots,e_k)\in\mathsf{Eq}^k$ from the $\mathsf{Gap3Lin}$ instance (X,Eq). For $e\in\mathsf{Eq}$ let $v_e\in\mathbb{F}_q^X$ denote the indicator vector on the three variables appearing in e.

It will be convenient to only keep the $U=(e_1,\ldots,e_k)$ that satisfy the following properties:

- The equations $e_1,\ldots,e_k$ are distinct and do not share variables.

- For any $i \neq j$ and pair of variables $x \in e_i$ and $y \in e_j$, the variables x and y do not appear together in any equation in the instance (X, Eq).

The fraction of $U = (e_1, \ldots, e_k)$ that do not satisfy the above is $O(k^2/n)$ which is negligible for us, and dropping them will only reduce our completeness by $o(1)$. This will not affect our analysis, and henceforth we will assume that all $U = (e_1, \ldots, e_k)$ satisfy the above properties. We now describe the 2-Prover-1-Round Games instance $\Psi = (\mathcal{A}, \mathcal{B}, E, \Sigma_1, \Sigma_2, \Phi)$. All vertices in the underlying graph will correspond to subspaces of $\mathbb{F}_q^X$.

4.1.1 The Vertices

For each question $U = (e_1, \ldots, e_k)$, let $H_U = \mathrm{span}(v_{e_1}, \ldots, v_{e_k})$, where v_{e_i} is the vector with ones at coordinates corresponding to variables appearing in e_i. We can think of the v_{e_i}'s as vectors from an underlying space $\mathbb{F}_q^U$. By the first property described above, $\dim(H_U) = k$ and $\dim(\mathbb{F}_q^U) = 3k$. The vertices of Ψ are:

$$\mathcal{A} = \{L \oplus H_U \mid U \in \mathcal{U}, L \subseteq \mathbb{F}_q^U, \dim(L) = 2\ell, L \cap H_U = \{0\}\},$$
$$\mathcal{B} = \{R \mid \exists U \in \mathcal{U}, \text{ s.t. } R \subseteq \mathbb{F}_q^U, \dim(R) = 2(1-\delta)\ell\}.$$

In words, the vertices on the side $\mathcal{A}$ are all 2ℓ-dimensional subspaces of some $\mathbb{F}_q^U$ for some $U \in \mathcal{U}$. For technical reasons, we require them to intersect H_U trivially (which is the case for a typical 2ℓ-dimensional space) and add to them the space H_U.[7] The vertices on the side $\mathcal{B}$ are all $2(1-\delta)\ell$ dimensional subspaces of $\mathbb{F}_q^U$.

4.1.2 The Alphabets

The alphabets Σ_1, Σ_2 have sizes $|\Sigma_1| = q^{2\ell}$ and $|\Sigma_2| = q^{2(1-\delta)\ell}$. For each vertex $L \oplus H_U \in \mathcal{A}$, let $\psi : H_U \to \mathbb{F}_q$ denote the function that satisfies the side conditions given by the equations in U. Namely, if $e_i \in U$ is the equation $\langle x, h_i \rangle = b_i$ for $x \in \mathbb{F}_q^U$, then $\psi(h_i) = b_i$. We say a linear function $f\colon L \oplus H_U \to \mathbb{F}_q$ satisfies the side conditions of U if $f|_{H_U} \equiv \psi$. In this language, for a vertex $L \oplus H_U$ we identify Σ_1 with

$$\{f : L \oplus H_U \to \mathbb{F}_q \mid f \text{ is linear function satisfying the side conditions of } U\}.$$

As $L \cap H_U = \{0\}$ and $\dim(L) = 2\ell$, it is easy to see that the above set indeed has size $q^{2\ell}$. For each right vertex R, we identify Σ_2 with

$$\{f : R \to \mathbb{F}_q \mid f \text{ is linear}\}.$$

4.1.3 The Edges

To define the edges, we first need the following relation on the vertices in A. Say that $(L \oplus H_U) \sim (L' \oplus H_{U'})$ if

$$L \oplus H_U \oplus H_{U'} = L' \oplus H_U \oplus H_{U'}.$$

Recall that all subspaces above are in $\mathbb{F}_q^X$ hence the direct sums and equality above are well defined. The relation described is in fact an equivalence relation and thus partitions the vertices in $\mathcal{A}$ into disjoint equivalence classes. It is clear that the relation is reflexive and symmetric, so we need only show that it is also transitive.

[7] This has the effect of collapsing L and L' such that $L \oplus H_U = L' \oplus H_U$ to a single vertex.

Lemma 4.1. *If $L_1 \oplus H_{U_1} \oplus H_{U_2} = L_2 \oplus H_{U_1} \oplus H_{U_2}$, and $L_2 \oplus H_{U_2} \oplus H_{U_3} = L_3 \oplus H_{U_2} \oplus H_{U_3}$, then*

$$L_1 \oplus H_{U_1} \oplus H_{U_3} = L_3 \oplus H_{U_1} \oplus H_{U_3}.$$

Proof. We "add" H_{U_1} to the second equation to obtain,

$$L_2 \oplus H_{U_1} \oplus H_{U_2} \oplus H_{U_3} = L_3 \oplus H_{U_1} \oplus H_{U_2} \oplus H_{U_3}.$$

Next, write $H_{U_2} = A \oplus B$, where A is the span of all vectors v_e for equations e in U_2 that are also in U_1 or U_3, while B is the span of all vectors v_e for equations $e \in U_2$ that are in neither U_1 nor U_3. It follows that $A \cap B = \{0\}$. Now note that any equation $e \in B$ has at most one variable that appears in an equation in U_1, and at most one variable that appears in an equation in U_2. Thus, each $e \in B$, has a "private variable", and as the equations in B are over disjoint sets of variables, this private variable does not appear in $U_1 \cup U_3 \cup (U_2 \setminus e)$. It follows that

$$B \cap L_1 \oplus L_3 \oplus H_{U_1} \oplus H_{U_3} = \{0\} \subset \mathbb{F}_q^X.$$

Indeed, by the above discussion any nonzero vector in $B \subseteq \mathbb{F}_q^X$ is nonzero on at least one coordinate of X (corresponding to a private variable), and no vector in $\mathbb{F}_q^{U_1}$ or $\mathbb{F}_q^{U_2}$ is supported on this coordinate.

Substituting $H_{U_2} = A \oplus B$ into the original equation yields,

$$L_1 \oplus (H_{U_1} \oplus H_{U_3} \oplus A) \oplus B = L_3 \oplus (H_{U_1} \oplus H_{U_3} \oplus A) \oplus B.$$

Since $A \subset H_{U_1} \oplus H_{U_3}$, this equivalent to

$$L_1 \oplus H_{U_1} \oplus H_{U_3} \oplus B = L_3 \oplus H_{U_1} \oplus H_{U_3} \oplus B.$$

As $B \cap L_1 \oplus L_3 \oplus H_{U_1} \oplus H_{U_3} = \{0\}$, it follows that

$$L_1 \oplus H_{U_1} \oplus H_{U_3} = L_3 \oplus H_{U_1} \oplus H_{U_3},$$

as desired. $\square$

By Lemma 4.1 the relation $\sim$ is indeed an equivalence relation and we may partition $\mathcal{A}$ into equivalence classes, $[L \oplus H_U]$. We call each class a clique and partition $\mathcal{A}$ into cliques:

$$\mathcal{A} = \mathsf{Clique}_1 \sqcup \cdots \sqcup \mathsf{Clique}_m.$$

The actual number of cliques, m, will not be important, but it is clear that such a number exists. The edges of our graph will be between vertices $L \oplus H_U$ and R if there exists $L' \oplus H_{U'} \in [L \oplus H_U]$ such that $L' \supseteq R$. The edges will be weighted according to a sampling process that we describe in the next section, which also explains the constraints on Ψ. For future reference, the following lemma will be helpful in defining the constraints:

Lemma 4.2. *Suppose $L \oplus H_U \sim L' \oplus H_{U'}$ and that $f : L \oplus H_U \to \mathbb{F}_q$ is a linear function satisfying the side conditions. Then there is a unique linear function $f' : L' \oplus H_{U'} \to \mathbb{F}_q$ that satisfies the side conditions such that there exists a linear function $g : L \oplus H_U \oplus H_{U'} \to \mathbb{F}_q$ satisfying the side conditions (of both U and U') such that*

$$g|_{L \oplus H_U} = f \quad \text{and} \quad g|_{L' \oplus H_{U'}} = f'.$$

In words, g is a linear extension of both f and f'.

Proof. Note that there is only one way to extend f to $L \oplus H_U \oplus H_{U'}$ in a manner that satisfies the side conditions given by U'. Let this function be g. We take f' to be $g|_{L' \oplus H_{U'}}$. $\square$

4.1.4 The Constraints

Suppose that T_1 is an assignment to $\mathcal{A}$ that assigns, to each vertex $L \oplus H_U$, a linear function $T_1[L \oplus H_U]$ satisfying the side conditions. Further suppose that T_2 is an assignment that assigns to each vertex $R \in \mathcal{B}$ a linear function on R. The verifier performs the following test, which also describes the constraints of Ψ:

1. Choose U uniformly at random from $\mathcal{U}$.
2. Choose $L \oplus H_U$ uniformly, where $\dim(L) = 2\ell$ and $L \cap H_U = \{0\}$, and choose $R \subseteq L$ of dimension $2(1-\delta)\ell$ uniformly.
3. Choose $L' \oplus H_{U'} \in [L \oplus H_U]$ uniformly
4. As in Lemma 4.2, extend $T_1[L' \oplus H_{U'}]$ to $L' \oplus H_{U'} \oplus H_U$ in the unique manner that respects the side conditions and let $\tilde{T}_1[L \oplus H_U]$ be the restriction of this extension to $L \oplus H_U$.
5. Accept if and only if $\tilde{T}_1[L \oplus H_U]|_R = T_2[R]$.

This finishes the description of our instance Ψ. It is clear that the running time and instance size is $n^{O(k)}$ and that the alphabet size is $O(q^{2\ell})$.

Before arguing about the completeness and soundness, we will present some necessary tools. As is usually the case, showing completeness is relatively easy, and all of the tools presented are for the much more complex soundness analysis.

5 Tools for Soundness Analysis

In this section we will present all of the tools needed to analyze the soundness of our PCP.

5.1 The 2ℓ versus $2\ell(1-\delta)$ subspace agreement test

We begin by discussing the 2ℓ versus $2\ell(1-\delta)$ test and our decoding theorem for it. In our setting, we have a question $U \in \mathcal{U}$ for the first prover, and we consider the 2ℓ versus $2\ell(1-\delta)$ test inside the space $\mathbb{F}_q^U$. In our setting this test passes with probability at least $\varepsilon \geqslant q^{-2\ell(1-\delta')}$ (where δ' is, say $\delta' = 1000\delta$) and we will want to use this fact to devise a strategy for the first prover. Below, we first state and prove a basic decoding theorem, and then deduce from it a quantitative better version that also incorporates the side conditions.

Let T_1 be a table that assigns, to each $L \in \mathsf{Grass}_q(\mathbb{F}_q^U, 2\ell)$, a linear function $T_1[L] : L \to \mathbb{F}_q$, and let T_2 be a table assigning to each $R \in \mathsf{Grass}_q(\mathbb{F}_q^U, 2(1-\delta)\ell)$ a linear function $T_2[R] \colon R \to \mathbb{F}_q$. We recall that $|U| = 3k \gg 2\ell$. In this section, we show that if tables T_1 and T_2 are ε-consistent, namely

$$\Pr_{\substack{L \in \mathsf{Grass}_q(\mathbb{F}_q^U, 2\ell) \\ R \in \mathsf{Grass}_q(\mathbb{F}_q^U, 2(1-\delta)\ell)}} [T_1[L]|_R = T_2[R] \mid R \subseteq L] \geqslant \varepsilon.$$

for $\varepsilon \geqslant q^{-2\ell(1-1000\delta)}$, then the table T_1 must have non-trivial agreement with a linear function on some zoom-in and zoom-out combination of constant dimension. To prove that, we use Lemma 2.3 along with an idea from [BKS19].

Theorem 5.1. *Suppose that tables T_1 and T_2 are ε-consistent where $\varepsilon \geqslant q^{-2\ell(1-1000\delta)}$. Then there exist subspaces $Q \subset W$ and a linear function $f : W \to \mathbb{F}_q$ such that:*

1. $\mathrm{codim}(Q) + \dim(W) = \frac{10}{\delta}$.

2. $f|_L \equiv T_1[L]$ *for* $\Omega(\varepsilon')$*-fraction of* 2ℓ*-dimensional* $L \in \mathsf{Zoom}[Q, W]$*,*

where $\varepsilon' = q^{-2\ell(1-1000\delta^2)}$.

Proof of Theorem 5.1. Consider the bipartite graph G whose sides are the vertices of $\mathsf{Grass}_q(n, 2\ell)$ and $\mathsf{Grass}_q(n, (1-\delta)2\ell)$, and its set of edges E consists of pairs (L, R) such that $L \supseteq R$. Consider the normalized adjacency operator $\mathcal{T} : L_2(\mathsf{Grass}_q(n, 2\ell)) \to L_2(\mathsf{Grass}_q(n, 2(1-\delta)\ell))$ of G, and let $\mathcal{T}^*$ be its adjoint operator.

Choose a linear function $f : \mathbb{F}_q^n \to \mathbb{F}_q$ uniformly at random and define the (random) sets of vertices

$$S_{L,f} = \{L \in \mathsf{Grass}_q(n, 2\ell) \mid f|_L \equiv T_1[L]\} \quad \text{and} \quad S_{R,f} = \{R \in \mathsf{Grass}_q(n, 2(1-\delta)\ell) \mid f|_R = T_2[R]\}.$$

Denote by $E(S_{L,f}, S_{R,f})$ the set of edges with endpoints in $S_{L,f}$ and $S_{R,f}$. We lower bound the expected size of $E(S_{L,f}, S_{R,f})$ over the choice of f. Note that for each edge $(L, R) \in E$ such that $T_1[L]|_R \equiv T_2[R]$, we have that $(L, R) \in E(S_{L,f}, S_{R,f})$ with probability $q^{-2\ell}$. Indeed, with probability $q^{-2\ell}$ we have that $T_1[L] \equiv f|_L$, and in that case we automatically get that $T_2[R] = T_1[L]|_R \equiv (f|_L)|_R = f|_R$. As the number of edges (L, R) such that $T_1[L]|_R \equiv T_2[R]$ is at least $\varepsilon|E|$, we conclude that

$$\mathop{\mathbb{E}}_{f}[|E(S_{L,f}, S_{R,f})|] \geqslant \varepsilon q^{-2\ell}|E|.$$

Note that we also have that

$$\mathop{\mathbb{E}}_{f}[\mu(S_{R,f})] = \mathop{\mathbb{E}}_{f}\left[\frac{|S_{R,f}|}{|R|}\right] = q^{-2\ell(1-\delta)}.$$

Using Linearity of Expectation, we get that

$$\mathop{\mathbb{E}}_{f}\left[|E(S_{L,f}, S_{R,f})| - \frac{1}{2}\varepsilon q^{2\delta\ell}\mu(S_{R,f})|E|\right] \geqslant \frac{1}{2}\varepsilon q^{-2\ell}|E|,$$

thus there exists f for which the random variable on the left hand side is at least $\frac{1}{2}\varepsilon q^{-2\ell}|E|$, and we fix f so that

$$|E(S_{L,f}, S_{R,f})| \geqslant \frac{1}{2}\varepsilon q^{2\delta\ell}\mu(S_{R,f})|E| + \frac{1}{2}\varepsilon q^{-2\ell}|E|. \tag{2}$$

We claim that $S_{L,f}$ is not (r, ε')-pseudo-random for $r = \frac{10}{\delta}$ and $\varepsilon' = q^{-2\ell(1-1000\delta^2)}$. Suppose for the sake of contradiction that this is not the case, and that $S_{L,f}$ is (r, ε')-pseudo-random. Denote $\alpha = \mu(S_{L,f})$ and $\beta = \mu(S_{R,f})$. By Lemma 2.3 for any $t \geqslant 4$ that is a power of 2 we have

$$\frac{1}{|E|}|E(S_L, S_R)| \leqslant q^{O_{t,r}(1)}\beta^{\frac{t-1}{t}}\varepsilon'^{\frac{t-1}{t}} + 5q^{-2r\delta\ell}\sqrt{\alpha\beta} \leqslant q^{O_{t,r}(1)}\beta^{\frac{t-1}{t}}\varepsilon'^{\frac{t-1}{t}}. \tag{3}$$

In the last inequality, we used the fact that by (2)

$$\beta|E| = |S_{R,f}|\frac{|E|}{|R|} \geqslant |E(S_{L,f}, S_{R,f})| \geqslant \frac{1}{2}\varepsilon q^{-2\ell}|E|,$$

so $\beta \geqslant \frac{1}{2}\varepsilon q^{-2\ell} \geqslant q^{-4\ell}$, and thus the second term on the middle of (3) is negligible compared to the first term there. Combining (2) and (3) gives us that

$$\frac{1}{2}\varepsilon q^{2\delta\ell}\beta \leqslant q^{O_{t,r}(1)}\beta^{\frac{t-1}{t}}\varepsilon'^{\frac{t-1}{t}}.$$

Simplifying, using the definition of ε', the fact that $\varepsilon \geqslant q^{-2\ell(1-1000\delta)}$ and the fact that $\beta \geqslant q^{-4\ell}$ we get

$$\frac{1}{2}q^{2\delta\ell} \leqslant q^{O_{t,r}(1)} q^{\frac{4\ell}{t}} q^{\left(\frac{2}{t}+2000\left(\delta^2\frac{t-1}{t}-\delta\right)\right)\ell}.$$

Investigating the second two exponents of q, we have that for $t \geqslant \frac{1}{\delta-\delta^2} \geqslant 2$,

$$\left(\frac{4}{t}+\frac{2}{t}+2000\left(\delta^2\frac{t-1}{t}-\delta\right)\right)\ell \leqslant -1994(\delta-\delta^2)\ell.$$

This implies that

$$\frac{1}{2}q^{2\delta\ell} \leqslant q^{O_\delta(1)} q^{1994\ell(\delta^2-\delta)} < 1,$$

and contradiction. It follows that $S_{L,f}$ is not (r,ε')-pseudo-random, and unraveling the definition of not being pseudo-random gives the conclusion of the theorem. □

5.1.1 Finding a Large Fraction of Successful Zoom-Ins

Theorem 5.1 asserts the existence of a good pair of zoom-in and zoom-out (Q,W) on which the table T_1 has good agreement with a global linear function. As discussed in the introduction, our argument requires a quantitatively version asserting that there is a good fraction of zoom-ins that work for us. Below, we state a strengthening of Theorem 5.1 achieving this; it easily follows from Theorem 5.1, and we defer the proof to Section B.

Theorem 5.2. *Suppose that tables T_1 and T_2 are ε-consistent for $\varepsilon \geqslant q^{-2\ell(1-1000\delta)}$. Then there exist positive integers r_1 and r_2 satisfying $r_1+r_2 = r = \frac{10}{\delta}$, such that for at least $q^{-5\ell^2}$-fraction of the r_1-dimensional subspaces Q, there exists a subspace $W \supseteq Q$ of codimension r_2 and a linear function $g_{Q,W}$ such that*

$$\Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell)}[g_{Q,W}|_L = T_1[L] \mid Q\subseteq L\subseteq W] \geqslant q^{-2\ell(1-1000\delta^2)}.$$

Proof. The proof is deferred to Section B. □

5.1.2 Incorporating Side Conditions for Zoom-Ins

Next, we require a version of Theorem 5.2 which also takes the side conditions into account.

Theorem 5.3. *Let U be a question to the first prover, let T_1 the first prover's table, including the side conditions, and suppose that*

$$\Pr_{\substack{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell),L\cap H_U=\{0\}\\ R\in\mathsf{Grass}_q(\mathbb{F}_q^U,2(1-\delta)\ell)}}[T_1[L\oplus H_U]|_R = T_2[R] \mid R\subseteq L] = \varepsilon \geqslant q^{-2(1-1000\delta)\ell}.$$

Then there are parameters r_1 and r_2 such that $r_1+r_2 \leqslant \frac{10}{\delta}$, such that for at least $q^{-6\ell^2}$ fraction of the r_1-dimensional subspaces $Q\subseteq\mathbb{F}_q^U$, there exists $W\subseteq\mathbb{F}_q^U$ of codimension r_2 containing $Q\oplus H_U$, and a global linear function $g_{Q,W}: W\to\mathbb{F}_q$ that respects the side conditions on H_U such that

$$\Pr_L[g_{Q,W}|_{L\oplus H_U} = T_1[L\oplus H_U] \mid Q\subseteq L\subseteq W] \geqslant \frac{q^{-2(1-1000\delta^2)\ell}}{5}.$$

Proof. For any $2k$-dimensional subspace A such that $H_U \oplus A = \mathbb{F}_q^U$, let T_A be the table given by $T_A[L] = T_1[L \oplus H_U]|_L$ for all 2ℓ-dimensional subspaces $L \subseteq A$. We can choose a 2ℓ-dimensional subspace L such that $L \cap H_U = \{0\}$ by first uniformly choosing A such that $H_U \oplus A = \mathbb{F}_q^U$, and then choosing $L \subseteq A$ of dimension 2ℓ uniformly. Thus, defining

$$p'(A) = \Pr_{L \subseteq A, R \subseteq L}[T_A[L]|_R = T_2[R]],$$

we have

$$\mathop{\mathbb{E}}_{A}[p'(A)] = \Pr_{L:\dim(L)=2\ell, L\cap H_U=\{0\}, R\subseteq L}[T_1[L \oplus H_U]|_R = T_2[R]] \geqslant \varepsilon.$$

In particular, for at least $\frac{\varepsilon}{4}$-fraction of A's, we have $p'(A) \geqslant \frac{\varepsilon}{4}$. For such A's, by Theorem 5.2, there exist positive integers r_1 and r_2 such that for at least $q^{-5\ell^2}$-fraction of r_1-dimensional zoom-ins Q, there exists a zoom-out $W' \supset Q$ of co-dimension r_2 and a linear function $g_{Q,W'}$ such that,

$$\Pr_{Q\subseteq L\subseteq W}[T_A[L] = g_{Q,W}|_L] \geqslant \frac{q^{-2(1-1000\delta^2)}}{4}.$$

Let $W = W' \oplus H_U$ and let $g_{Q,W} : W \to \mathbb{F}_q$ be the unique extension of $g_{Q,W'}$ to W satisfying the side conditions. We claim that

$$\Pr_{L:L\cap H_U=\{0\}}[g_{Q,W}|_{L\oplus H_U} = T_1[L \oplus H_U] \mid Q \subseteq L \subseteq W] \geqslant \frac{q^{-2(1-1000\delta^2)}}{5}.$$

Indeed, for each $Q \subseteq L' \subseteq W'$ there are an equal number of L such that $Q \subseteq L \subseteq W$ and $L' \oplus H_U = L \oplus H_U$, so

$$\begin{aligned}
&\Pr_{L:L\cap H_U=\{0\}}[g_{Q,W}|_{L\oplus H_U} = T_1[L \oplus H_U] \mid Q \subseteq L \subseteq W] \\
&\geqslant \Pr_{Q\subseteq L\subseteq W}[T_A[L] = g_{Q,W}|_L] - \Pr_{Q\subseteq L\subseteq W}[L \cap H_U \neq \{0\}] \\
&\geqslant \frac{q^{-2(1-1000\delta^2)}}{5}.
\end{aligned}$$

To conclude, we see that sampling A and then $Q \subseteq A$ of dimension r_1, we get that there is a zoom-out W and a function $g_{Q,W}$ satisfying the conditions in the theorem with probability at least $\frac{\varepsilon}{4}q^{-5\ell^2}$. As the marginal distribution over Q is $q^{-\Omega(k)}$-close to uniform over all r_1-dimensional subspaces the conclusion follows. $\square$

5.2 The Covering Property

In this section, we present the so called "covering property", which is a feature of our PCP construction that allows us to move between the first prover's distribution over 2ℓ-dimensional subspaces of $\mathbb{F}_q^U$ and the second prover's distribution over 2ℓ-dimensional subspaces of $\mathbb{F}_q^V$. Similar covering properties are shown in [KS13, KMS17]; however, obtaining the optimal quadratic-programming hardness result in Theorem 1.4 requires a stronger analysis that goes beyond the covering properties of [KS13, KMS17]. We are able to obtain a covering property with the following parameters:

$$k = q^{2(1+c)\ell} \quad , \quad \beta = q^{-2(1+2c/3)\ell}, \tag{4}$$

where $c > 0$ is a constant arbitrarily small relative to δ.

5.2.1 The Basic Covering Property

To start, we state a basic form of the improved covering property that is required in our analysis and defer its proof to Appendix C. Fix a question $U = (e_1, \ldots, e_k)$ to the first prover and recall that $H_U = \mathrm{span}(x_{e_1}, \ldots, x_{e_k})$ where x_{e_i} is the vector that is one at coordinates corresponding to variables in e_i and 0 elsewhere. The covering property we show will relate the following two distributions:
$\mathcal{D}$:

- Choose $x_1, \ldots, x_{2\ell} \in \mathbb{F}_q^U$ uniformly.
- Output the list $(x_1, \ldots, x_{2\ell})$.

$\mathcal{D}'$:

- Choose $V \subseteq U$ according to the Outer PCP.
- Choose $x'_1, \ldots, x'_{2\ell} \in \mathbb{F}_q^V$ uniformly, and lift these vectors to $\mathbb{F}_q^U$ by inserting 0's into the missing coordinates.
- Choose $w_1, \ldots, w_{2\ell} \in H_U$ uniformly, and set $x_i = x'_i + w_i$ for $1 \leqslant i \leqslant 2\ell$.
- Output the list $(x_1, \ldots, x_{2\ell})$.

With these two definitions, the covering property used in prior works asserted that the distribution $\mathcal{D}$ is statistically close to a variant of the distribution $\mathcal{D}'$. This closeness is not good enough for us, as we will want to consider events of rather small probability under $\mathcal{D}$ and still assert that their probability is roughly the same in $\mathcal{D}'$. First, in these earlier works, the distribution $\mathcal{D}'$ was generated by a similar process to the above without the addition of the random vectors $w_1, \ldots, w_{2\ell}$ from H_U. As explained in the introduction however, this distribution is not good enough for the purpose of Theorem 1.4, and we must consider the distribution $\mathcal{D}'$ above. Second, the notion of statistical closeness is too rough for us, and we show that in fact, almost all inputs x are assigned the same probability under these two distributions up to factor $1 + o(1)$.

More precisely, set $\eta = q^{-100\ell^{100}}$ throughout this subsection. Our covering property is the following statement:

Lemma 5.4. *Let η be a parameter such that $q^{-100\ell^{100}} \leqslant \eta \leqslant 1/2$. There exists a small set $E \subseteq \left(\mathbb{F}_q^U\right)^{2\ell}$ such that both $\mathcal{D}(E)$ and $\mathcal{D}'(E)$ are at most η^{40}, and for all $(x_1, \ldots, x_{2\ell}) \notin E$ we have*

$$0.9 \leqslant \frac{D(x_1, \ldots, x_{2\ell})}{D'(x_1, \ldots, x_{2\ell})} \leqslant 1.1.$$

Proof. The proof is deferred to Appendix C.1. □

5.2.2 The Covering Property with Zoom-ins

Lemma 5.4 represents the most basic form of the covering property, and for out application we require a version of it that incorporates zoom-ins and advice. Namely, we will actually interested in the case where $\mathcal{D}$ and $\mathcal{D}'$ are conditioned on some r_1-dimensional zoom-in Q, for an arbitrary dimension $r_1 \leqslant \frac{10}{\delta}$. To make notation simpler, let us write $x = (x_1, \ldots, x_{2\ell})$ and use $\mathrm{span}_{r_1}(x)$ to denote $\mathrm{span}(x_1, \ldots, x_{r_1})$. Additionally, define

$$\mathcal{D}_Q(\mathcal{L}) = \mathcal{D}(x \in \mathcal{L} \mid \mathrm{span}_{r_1}(x) = Q) = \frac{\mathcal{D}(\{x \in \mathcal{L} \mid \mathrm{span}_{r_1}(x) = Q\})}{\mathcal{D}(\{x \mid \mathrm{span}_{r_1}(x) = Q\})}.$$

From Lemma 5.4 we can conclude that for any $\mathcal{L} \subseteq (\mathbb{F}_q^U)^{2\ell}$ that is not too small, the measure $\mathcal{D}'_Q(\mathcal{L})$ is within at least a constant factor of $\mathcal{D}_Q(\mathcal{L})$ for nearly all Q.

Lemma 5.5. *For any $\mathcal{L} \subseteq (\mathbb{F}_q^U)^{2\ell}$, we have*

$$\Pr_Q\left[\mathcal{D}'_Q(\mathcal{L}) \geqslant 0.8 \cdot \mathcal{D}_Q(\mathcal{L}) - \eta^{20}\right] \geqslant 1 - 2\eta^{20},$$

where Q is the span of r_1 uniformly random vectors in $\mathbb{F}_q^U$.

Proof. Throughout the proof all of the expectations and probabilities over Q choose Q as in the lemma statement. Let E be the small set of points from Lemma 5.4. By assumption we have

$$\mathop{\mathbb{E}}_Q[\mathcal{D}_Q(E)] \leqslant \eta^{40} \quad \text{and} \quad \mathop{\mathbb{E}}_Q[\mathcal{D}'_Q(E)] \leqslant \eta^{40}.$$

Thus, by Markov's inequality, we have that with probability at least $1 - 2\eta^{20}$, we have $\mathcal{D}_Q(E), \mathcal{D}'_Q(E) \leqslant \eta^{20}$. In this case we have,

$$\begin{aligned}\sum_{\mathrm{span}_{r_1}(x)=Q} D(x) \geqslant \sum_{\mathrm{span}_{r_1}(x)=Q, x\in\overline{E}} D(x) &\geqslant \sum_{\mathrm{span}_{r_1}(x)=Q, x\in\overline{E}} 0.9 \cdot \mathcal{D}'(x) \\ &= 0.9 \cdot \left(\sum_{\mathrm{span}_{r_1}(x)=Q} \mathcal{D}'(x) - \sum_{\mathrm{span}_{r_1}(x)=Q, x\in E} \mathcal{D}'(x)\right),\end{aligned}$$

where we applied Lemma 5.4 in the second transition. Dividing both sides by $\sum_{\mathrm{span}_{r_1}(x)=Q} \mathcal{D}'(x)$ gives that

$$\frac{\sum_{\mathrm{span}_{r_1}(x)=Q} D(x)}{\sum_{\mathrm{span}_{r_1}(x)=Q} \mathcal{D}'(x)} \geqslant 0.9(1 - \mathcal{D}'_Q(E)) \geqslant 0.9(1 - \eta^{20}) \geqslant 0.89. \tag{5}$$

It follows that

$$\begin{aligned}\mathcal{D}'_Q(\mathcal{L}) &= \frac{\sum_{x\in\mathcal{L}, \mathrm{span}_{r_1}(x)=Q} \mathcal{D}'(x)}{\sum_{\mathrm{span}_{r_1}(x)=Q} \mathcal{D}'(x)} \\ &\geqslant \frac{0.89 \sum_{x\in\mathcal{L}\cap\overline{E}, \mathrm{span}_{r_1}(x)=Q} \mathcal{D}'(x)}{\sum_{\mathrm{span}_{r_1}(x)=Q} \mathcal{D}(x)} \\ &\geqslant \frac{0.9 \cdot 0.89 \cdot \sum_{x\in\mathcal{L}\cap\overline{E}, \mathrm{span}_{r_1}(x)=Q} \mathcal{D}(x)}{\sum_{\mathrm{span}_{r_1}(x)=Q} \mathcal{D}(x)} \\ &\geqslant 0.8 \cdot \mathcal{D}_Q(L) - \frac{\sum_{\mathrm{span}_{r_1}(x)=Q, x\in E} \mathcal{D}(x)}{\sum_{\mathrm{span}_{r_1}(x)=Q} \mathcal{D}(x)} \\ &= 0.8 \cdot \mathcal{D}_Q(L) - \mathcal{D}_Q(E) \\ &\geqslant 0.8 \cdot \mathcal{D}_Q(L) - \eta^{20},\end{aligned}$$

where we apply Equation (5) in the second transition and the assumption $\mathcal{D}_Q(E) \leqslant \eta^{20}$ in the last transition. $\square$

5.2.3 The Covering Property for the Advice

We will also need a similar, and simpler, version of the above lemma that applies to r_1-dimensional subspaces for some constant $r_1 = O(\delta^{-1})$. This is to handle the fact that the zoom-in Q is sampled uniformly from $\mathbb{F}_q^V$ after V is chosen according to the outer PCP, and then lifted to a subspace over $\mathbb{F}_q^U$, instead of uniformly from $\mathbb{F}_q^U$. Formally, let $\mathcal{D}'_{r_1}$ denote the former distribution over r_1-dimensional subspaces $Q \subseteq \mathbb{F}_q^U$ and let $\mathcal{D}_{r_1}$ denote the latter distribution over r_1-dimensional subspaces $Q \subseteq \mathbb{F}_q^U$. These are the same as the distributions $\mathcal{D}$ and $\mathcal{D}'$ of the previous subsection except over $\left(\mathbb{F}_q^U\right)^{r_1}$ instead of $\left(\mathbb{F}_q^U\right)^{2\ell}$. We show the following.

Lemma 5.6. *Let $\mathcal{Q}$ be a set of r_1-dimensional subspaces in $\mathbb{F}_q^U$ satisfying $\mathcal{D}_{r_1}(\mathcal{Q}) \geqslant q^{-10\ell^{10}}$. Then,*

$$\mathcal{D}'_{r_1}(\mathcal{Q}) \geqslant 0.8 \cdot \mathcal{D}_{r_1}(Q).$$

Proof of Lemma 5.6. Take E from Lemma 5.4, and define

$$\mathcal{L} = \{(x_1, \dots, x_{2\ell}) \in \left(\mathbb{F}_q^U\right)^{2\ell} \mid \operatorname{span}(x_1, \dots, x_{r_1}) \in \mathcal{Q}\}.$$

Then $\mathcal{D}_{r_1}(\mathcal{Q}) = \mathcal{D}(\mathcal{L})$, $\mathcal{D}'_{r_1}(\mathcal{Q}) = \mathcal{D}'(\mathcal{L})$ and

$$\mathcal{D}'_{r_1}(\mathcal{Q}) = \mathcal{D}'(\mathcal{L}) \geqslant \sum_{\substack{x \in \mathcal{L} \\ x \notin E}} \mathcal{D}'(x) \geqslant 0.9 \cdot \sum_{\substack{x \in \mathcal{L} \\ x \notin E}} \mathcal{D}(x) \geqslant 0.9 \cdot \mathcal{D}(\mathcal{L}) - \mathcal{D}(E) \geqslant 0.8 \cdot \mathcal{D}(\mathcal{L}) = 0.8 \cdot \mathcal{D}_{r_1}(\mathcal{L}).$$

where we use Lemma 5.4 in the second transition, and the fact that $\mathcal{D}(E) \leqslant \eta$ and $\mathcal{D}(\mathcal{L}) = \mathcal{D}_{r_1}(\mathcal{Q}) \geqslant q^{-10\ell^{10}}$ in the penultimate transition. □

5.3 The Number of Maximal Zoom-Outs is Bounded

In Theorem 5.2, we showed that the two provers can agree on a zoom-in with reasonable probability using their advice. The same cannot be said for zoom-outs however, and to circumvent this issue we must develop further tools. In this section, we define the notion of *maximal* zoom-outs and show that for a fixed zoom-in Q, the number of maximal zoom-outs is bounded.

5.3.1 Generic Sets of Subspaces

One of our primary concerns with respect to zoom-outs is that it is possible for a prover to have many good zoom-outs to choose from (so that independent sampling doesn't work) but not enough to allow for advice-type solution. To deal with large collections of zoom-outs we define a special property of zoom-outs that is called "genericness". To motivate it, note that if $W_1, W_2 \subseteq V$ are distinct subspaces of co-dimension r, then $W_1 \cap W_2$ is a subspace whose co-dimension is between $2r$ and $r+1$. For a typical pair of subspaces the intersection $W_1 \cap W_2$ has dimension $2r$, in which case we say they are generic. Genericness is useful probabilistically, since if W_1, W_2 are generic then the event that a randomly chosen 2ℓ-dimensional subspace is contained in W_1, and the event it is contained in W_2, are almost independent. Below is a more general and formal definition:

Definition 5.7. *We say that a set $\mathcal{S} = \{W_1, \dots, W_N\}$ of codimension r subspaces of V is t-generic with respect to V if for any t-distinct subspaces, say $W_{i_1}, \dots, W_{i_t} \in \mathcal{S}$, we have $\operatorname{codim}(\bigcap_{1 \leqslant j \leqslant t} W_{i_j}) = t \cdot r$. When the ambient space V is clear from context we simply say that $\mathcal{S}$ is t-generic.*

We remark that any set of subspaces that is t-generic with respect to V is also t'-generic with respect to V for any $t' < t$. In this section, we will show a couple of results regarding generic sets of subspaces that will be used to bound the number of maximal zoom-outs in Section 8.

The result we need is a sunflower-type lemma, stating that any large set of codimension r subspaces inside V contains a large set of subspaces that are t-generic with respect to V' for some $V' \subseteq V$. Below is a formal statement.

Lemma 5.8. *Let $t, r \in \mathbb{N}$ be integers and let $\mathcal{S} = \{W_1, \ldots, W_N\}$ be a set of N subspaces of co-dimension r inside of V. Then there exists a subspace $V' \subseteq V$ and a set of subspaces $\mathcal{S}' \subseteq \mathcal{S}$ such that:*

- $|\mathcal{S}'| \geqslant \frac{N^{\frac{1}{(r+1)\cdot(t-1)!}}}{q^r}$.
- *Each $W_i \in \mathcal{S}'$ is contained in V' and has co-dimension s with respect to V', where $s \leqslant r$.*
- *$\mathcal{S}'$ is t-generic with respect to V'.*

In order to show Lemma 5.8, we introduce two necessary lemmas. The first, Lemma 5.9, states that for $j \geqslant 2$, any j-generic set of subspaces contains a large $(j+1)$-generic set of subspaces. The second, Lemma 5.10,states that either a set of subspaces is already 2-generic, or there are many subspaces in the set that are contained in the same hyperplane. Using this lemma, we can start from a large set of subspaces $\mathcal{S}$ inside of an ambient space V and iteratively reduce to the dimension of the ambient space until we find a 2-generic set of subspaces relative to the ambient space. Indeed, either the set $\mathcal{S}$ is already 2-generic, or there is a hyperplane $V' \subseteq V$ such that the set of subspaces in $\mathcal{S}$ contained in V' is large. Taking this set to be the new $\mathcal{S}$ and V' to be the new ambient space, we obtain a, still, large set of subspaces, whose codimension is now one less. We may repeat this process until $\mathcal{S}$ is a set of hyperplanes in the ambient space, at which point it will be 2-generic.

Lemma 5.9. *Let $\mathcal{S} = \{W_1, \ldots, \mathcal{W}_N\}$ be a set of N-subspaces of codimension r inside of V that is j-generic with respect to V, then there is a subset $\{W_1, \ldots, W_{N'}\} \subseteq \mathcal{S}$ of size $N' \geqslant \frac{N^{1/j}}{q^r}$ that is $(j+1)$-generic with respect to V.*

Proof. Fix any j distinct subspaces in $\mathcal{S}$, say $W_1, \ldots, W_j$ and let $W = W_1 \cap \cdots \cap W_j$. Since $\mathcal{S}$ is j-generic, $\operatorname{codim}(W) = j \cdot r$. We claim that there are at most $q^{j \cdot r}$ subspaces $W_{i'} \in \mathcal{S} \setminus \{W_1, \ldots, W_j\}$ such that $\operatorname{codim}(W \cap W_{i'}) \leqslant (j+1)r - 1$. Call such subspaces bad and suppose for the sake of contradiction that there are greater than q^{jr} bad subspaces $W_{i'}$. Then for each bad $W_{i'}$ we have,

$$\begin{aligned}\dim(W_{i'} \oplus W) &= \dim(W_{i'}) + \dim(W) - \dim(W_{i'} \cap W) \\ &\leqslant (\dim(V) - r) + (\dim(V) - jr) - (\dim(V) - (j+1)r - 1) \\ &= \dim(V) - 1.\end{aligned}$$

Therefore, for each $W_{i'}$, the space $W \oplus W_{i'}$ is contained in a hyperplane H such that $H \supseteq W$. There are at most $q^{\operatorname{codim}(W)} - 1 = q^{j \cdot r} - 1$ hyperplanes H containing W, and by the pigenhole principle it follows that there are two bad subspaces say $W_{i'_1}, W_{i'_2}$ that are both contained in the same hyperplane H. This is a contradiction however, as by the j-genericness of $\mathcal{S}$, we must have

$$\begin{aligned}\dim(W_{i'_1} \oplus W_{i'_2}) &= \dim(W_{i'_1}) + \dim(W_{i'_2}) - \dim(W_{i'_1} \cap W_{i'_2}) \\ &= 2(\dim(V) - r) - (\dim(V) - 2r) \\ &= \dim(V),\end{aligned}$$

and hence $W_{i'_1}$ and $W_{i'_2}$ cannot both be contained in the hyperplane H.

The lemma now follows from the claim we have just shown. Construct a subset $\mathcal{S}'$ greedily as follows:

1. Initialize $\mathcal{S}'$ by picking j arbitrary subspaces from $\mathcal{S}$ and inserting them to $\mathcal{S}'$.

2. For any j subspaces in $\mathcal{S}'$, say $W_1, \ldots, W_j$, remove any $W' \in \mathcal{S}$ which is bad for them.

3. If $\mathcal{S}$ is not empty, pick some $W \in \mathcal{S}$, insert it to $\mathcal{S}'$ and iterate.

Note that trivially, the collection $\mathcal{S}'$ will be $(j+1)$-generic in the end of the process. To lower bound the size of $\mathcal{S}'$, note that when $|\mathcal{S}'| = s$, the number of elements from $\mathcal{S}$ that have been deleted is at most $s^j q^{jr}$, and hence so long as this value is at most N, we may do another iteration. Thus, we must have that $s \geqslant \left(\frac{N}{q^{jr}}\right)^{1/j} = \frac{N^{1/j}}{q^r}$ when the process terminates. □

Lemma 5.10. *Let $\{W_1, \ldots, W_N\}$ be a set of subspaces of V of codimension r. Then for any integer $m \geqslant 1$, at least one of the following holds.*

- *There are m subspaces, way $W_1, \ldots, W_m$ such that for every pair $1 \leqslant i \neq j \leqslant m$, $\operatorname{codim}(W_i \cap W_j) = 2r$.*

- *There is a subspace $V' \subseteq V$ of co-dimension 1 that contains $N' = \frac{N}{mq^r}$ of these subspaces, say $W_1, \ldots, W_{N'}$.*

Proof. Note that for any $1 \leqslant i \neq j \leqslant N$, we have $\operatorname{codim}(W_i \cap W_j) \leqslant 2r$. Consider the graph with vertices $W_1, \ldots, W_N$ with (W_i, W_j) an edge if and only if $i \neq j$ and $\operatorname{codim}(W_i \cap W_j) \leqslant 2r-1$. If every vertex in this graph has degree at most $\frac{N}{m}$, then we are done as there is an independent set of size m and these subspaces satisfy the first condition. Suppose this is not the case. Then there is a vertex, say W_N, that has $\frac{N}{m}$ neighbors, say $W_1, \ldots, W_{\frac{N}{m}}$. For $1 \leqslant i \leqslant \frac{N}{m}$, we have $\operatorname{codim}(W_N \cap W_i) \leqslant 2r-1$, so

$$\dim(W_i \oplus W_N) = \dim(W_i) + \dim(W_N) - \dim(W_i \cap W_N) \leqslant \dim(V) - 1.$$

Thus $W_i \oplus W_N$ is always contained in a codimension 1 subspace of V that contains W_N. Since the number of such subspaces is $q^r - 1$, there must exist one subspace, say V', that contains at least $\frac{N}{mq^r}$ of the subspaces in the list $W_1, \ldots, W_{\frac{N}{m}}$. □

Repeatedly applying Lemma 5.10 yields the following corollary.

Corollary 5.11. *Let $\{W_1, \ldots, W_N\}$ be a set of subspaces of V of codimension r with respect to V. There exists a subspace $V' \subset V$, an integer $1 \leqslant s \leqslant r$, and a subset of $m \geqslant \frac{N^{\frac{1}{r+1}}}{q^r}$, say $\{W_1, \ldots, W_m\}$, all contained in V' such that,*

- *Each W_i, $1 \leqslant i \leqslant m$, has codimension s with respect to V'.*

- *Each $W_i \cap W_j$, $1 \leqslant i \neq j \leqslant m$, has codimension $2s$.*

Proof. To start set $V' = V$. If the W_i's have codimension 1 in V' then the result holds.

Otherwise, if the conclusion does not hold, then apply Lemma 5.10 with $m = \frac{N^{\frac{1}{r+1}}}{q^r}$. Either the first condition of Lemma 5.10 holds and we are done, or we can find a new subspace, V'', of codimension 1 inside the current V' containing at least $\frac{N}{mq^r}$ of the W_i's. Set $V' = V''$ and repeat. Note that the codimension of the W_i's with respect to V' drops by 1 after every iteration, so we will repeat at most r times before reaching the desired conclusion. This yields a list of W_i's that satisfy the conditions of size at least

$$\frac{N}{(mq^r)^r} = N^{\frac{1}{r+1}} \geqslant m. \qquad \square$$

With Corollary 5.11 and Lemma 5.9, we can prove Lemma 5.8

Proof of Lemma 5.8. By Corollary 5.11, there is a set $\mathcal{S}'$ of size $|\mathcal{S}'| \geqslant \frac{N^{\frac{1}{r+1}}}{q^r}$ and $V' \subseteq V$ such that $\mathcal{S}'$ is 2-generic with respect to V', and each W_i has (the same) codimension $s \leqslant r$ with respect to V'. Applying Lemma 5.9 $t-2$ times, there is a set of t-generic subspaces relative to V', $\mathcal{S}'' \subseteq \mathcal{S}'$, of size

$$\left|\mathcal{S}''\right| \geqslant \left(\left(\left(\left(\frac{|\mathcal{S}'|}{q^r}\right)^{\frac{1}{2}} \cdot \frac{1}{q^r}\right)^{\frac{1}{3}} \cdots\right) \cdot \frac{1}{q^r}\right)^{\frac{1}{t-1}} \geqslant \frac{|\mathcal{S}'|^{\frac{1}{(t-1)!}}}{q^r} \geqslant \frac{N^{\frac{1}{(r+1)\cdot(t-1)!}}}{q^r}. \qquad \square$$

In addition to Lemma 5.8, we state another useful feature of generic sets of subspaces, formalized in Lemma 5.12 below. The lemma asserts that if a collection $\{W_1, \ldots, W_N\}$ is generic, and one zooms-outs from the ambient space V into a hyperplane H, then one gets an induced collection $\{W_1 \cap H, \ldots, W_N \cap H\}$ which is almost as generic.

Lemma 5.12. *Let $\mathcal{S} = \{W_1, \ldots, W_N\}$ be a set of subspaces of codimension r that is t-generic with respect to some space V for an even integer t, and let H be a hyperplane in V. Then the set of subspaces $\mathcal{S}' = \{W_1 \cap H, \ldots, W_N \cap H\}$ can be made a $\frac{t}{2}$-generic set of subspaces with respect to H with codimension r inside of H by removing at most $\frac{t}{2}$ subspaces $W_i \cap H$ from it.*

Proof. Suppose that $\mathcal{S}'$ is not $\frac{t}{2}$-generic with respect to H with codimension r inside of H, as otherwise we are done. In this case, there must exist $\frac{t}{2}$ distinct subspaces, say $W_1 \cap H, \ldots, W_{\frac{t}{2}} \cap H \in \mathcal{S}'$ such that

$$\operatorname{codim}\left(W_1 \cap \cdots \cap W_{\frac{t}{2}} \cap H\right) < \frac{t}{2} \cdot r + 1,$$

where the codimension is with respect to V. However, since $\mathcal{S}$ is t-generic (and thus $\frac{t}{2}$-generic as well) with respect to V, this implies that

$$W_1 \cap \cdots \cap W_{\frac{t}{2}} \subseteq H.$$

Now delete $W_1 \cap H, \ldots, W_{\frac{t}{2}} \cap H$ from $\mathcal{S}'$. We claim that the resulting set is $\frac{t}{2}$-generic with respect to H. Suppose for the sake of contradiction that it is not. Then there must be another $\frac{t}{2}$ distinct subspaces, say $W_{\frac{t}{2}+1} \cap H, \ldots, W_t \cap H \in \mathcal{S}'$ such that

$$W_{\frac{t}{2}+1} \cap \cdots \cap W_t \subseteq H.$$

This would imply

$$\left(W_1 \cap \cdots \cap W_{\frac{t}{2}}\right) \oplus \left(W_{\frac{t}{2}+1} \cap \cdots \cap W_t\right) \subseteq H.$$

This is a contradiction however, as $\mathcal{S}$ is t-generic with respect to V, so $\operatorname{codim}(W_1 \cap \cdots \cap W_t) = tr$, and

$$\begin{aligned}
&\dim\left(W_1 \cap \cdots \cap W_{\frac{t}{2}} \oplus W_{\frac{t}{2}+1} \cap \cdots \cap W_t\right) \\
&\qquad = \dim\left(W_1 \cap \cdots \cap W_{\frac{t}{2}}\right) + \dim\left(W_{\frac{t}{2}} \oplus W_{\frac{t}{2}+1} \cap \cdots \cap W_t\right) - \dim\left(W_1 \oplus W_2 \cap \cdots \cap W_t\right) \\
&\qquad = 2\left(\dim(V) - \frac{t}{2}r\right) - \dim(V) + tr = \dim(V) > \dim(H),
\end{aligned}$$

and contradiction. $\square$

Lemma 5.13. *Let $\mathcal{W}$ be a set of subspaces that is 2^K generic with respect to V and let B be a subspace of codimension j. Then, the set of subspaces,*

$$\mathcal{W}_B = \{W_i \cap B \mid W_i \in \mathcal{W}\},$$

can be made 2^{K-j} generic with respect to B by removing at most $j2^{K-1}$ subspaces.

Proof. There is a sequence of subspaces $V = B_0 \supseteq B_1 \supseteq \cdots \supseteq B_j = B$, such that B_{i+1} is a hyperplane inside of B_i. Do the following,

1. Initialize $\mathcal{W}_0 = \mathcal{W}$ and set $i = 1$.
2. Set $\mathcal{W}_i = \{W_k \cap B_i \mid W_k \in \mathcal{W}_{i-1}\}$, and then remove the minimal number of subspaces to turn $\mathcal{W}_i$ into a 2^{K-i}-generic collection with respect to B_i.
3. Stop if $i = j$, otherwise, increase i by 1 and return to step 2.

It is clear that the output is a set of subspaces $\mathcal{W}_j \subseteq \mathcal{W}_B$ that is 2^{K-j}-generic with respect to B. Furthermore, during each iteration, at most $2^{K-i-1} \leqslant 2^{K-1}$ subspaces are removed by Lemma 5.12, and the result follows. $\square$

5.3.2 The Sampling Lemma

As explained earlier, the notion of genericness is useful probabilistically, and in this section we state and prove a sampling lemma about generic collections which is necessary for our analysis. Fix an arbitrary zoom-in $Q \subseteq V$ of dimension a, and let $\mathcal{S} = \{W_1, \dots, W_m\}$ be a 2-generic collection of subspaces of V of codimension r all containing Q. Also let $\mathcal{A}$ be a set of j-dimensional subspaces containing Q. For the remainder of this subsection, use $\mathsf{Zoom}[Q, V]$ to denote the set of j-dimensional subspaces in V containing Q. Consider the following two probability measures over $\mathsf{Zoom}[Q, V]$:

1. The distribution μ which is uniform over $\mathsf{Zoom}[Q, V]$.
2. The distribution ν, wherein a subspace is sampled by first picking $i \in \{1, \dots, m\}$ uniformly and then sampling a subspace from $\mathsf{Zoom}[Q, W_i]$ uniformly.

The main content of this section is the following lemma, asserting that the measures μ and ν are close in statistical distance provided that m is large. More precisely:

Lemma 5.14. *For any $\mathcal{L} \subset \mathsf{Zoom}[Q, V]$ we have*

$$|\nu(\mathcal{L}) - \mu(\mathcal{L})| \leqslant \frac{3q^{\frac{r}{2}(j-a)}}{\sqrt{m}}.$$

We now set up some notations for the proof of Lemma 5.14. For $L \in \mathsf{Zoom}[Q, V]$ let

$$N(L) = |\{W_i \in \mathcal{S} \mid L \subseteq W_i\}|,$$

and for an arbitrary pair of distinct $W_i, W_{i'} \in \mathcal{S}$ define the following quantities:

$$D = |\{L \in \mathsf{Zoom}[Q,V] \mid L \subseteq W_i\} = \begin{bmatrix} n-a-r \\ j-a \end{bmatrix}_q, \qquad p_1 = \Pr_{L \in \mathsf{Zoom}[Q,V]}[L \subseteq W_i]$$

$$p_2 = \Pr_{L \in \mathsf{Zoom}[Q,V]}[L \subseteq W_i \cap W_{i'}] = \frac{\begin{bmatrix} n-a-2r \\ j-a \end{bmatrix}_q}{\begin{bmatrix} n-a \\ j-a \end{bmatrix}_q}. \tag{6}$$

We note that all of these quantities are well defined as they do not depend on the identity of W_i and $W_{i'}$. The first two equations are clear and the third uses the fact that $\mathcal{S}$ is 2-generic. Thus,

$$\mu(L) = \frac{1}{|\mathsf{Zoom}[Q,V]|} = \frac{p_1}{D}, \qquad \nu(L) = \frac{N(L)}{m} \cdot \frac{1}{D} = \frac{N(L)}{mD}. \tag{7}$$

For the first equation, we are using the fact that $|\mathsf{Zoom}[Q,V]| \cdot p_1 = D$, while the second equation is evident. In the following claim we analyze the expectation and variance of $N(L)$ when L is chosen uniformly form $\mathsf{Zoom}[Q,V]$:

Claim 5.15. $\mathbb{E}_{L \in \mathsf{Zoom}[Q,V]}[N(L)] = p_1 m$ *and* $\mathrm{var}(N(L)) \leqslant p_1 m$. *where the variance is over uniform* $L \in \mathsf{Zoom}[Q,V]$.

Proof. By linearity of expectation

$$\mathbb{E}_{L \in \mathsf{Zoom}[Q,V]}[N(L)] = \sum_{i=1}^{m} \Pr_{L \in \mathsf{Zoom}[Q,V]}[L \subseteq W_i] = p_1 m,$$

and we move on to the variance analysis. To bound $\mathbb{E}_{L \in \mathsf{Zoom}[Q,V]}[N(L)^2]$, write

$$\begin{aligned}
\mathbb{E}_{L \in \mathsf{Zoom}[Q,V]}[N(L)^2] &= \mathbb{E}_{L \in \mathsf{Zoom}[Q,V]}\left[\left(\sum_{i=1}^{m} \mathbb{1}_{L \subseteq W_i}\right)^2\right] \\
&\leqslant m \cdot \Pr_{L \in \mathsf{Zoom}[Q,V]}[L \subseteq W_i] + m^2 \cdot \Pr_{L \in \mathsf{Zoom}[Q,V]}[L \subseteq W_i \cap W_{i'}] \\
&= p_1 m + p_2 m^2.
\end{aligned}$$

It follows that,

$$\mathrm{var}(N(L)) \leqslant p_1 m + p_2 m^2 - p_1^2 m^2.$$

Finally note that p_2 and p_1^2 are nearly the same value, and it can be checked using Equation (6) that $p_2 \leqslant p_1^2$ and so $\mathrm{var}(N(U)) \leqslant p_1 m$. $\square$

Combining Chebyshev's inequality with Claim 5.15, we conclude the following lemma which will be useful for us later on.

Lemma 5.16. *For any $c > 0$ it holds that*

$$\Pr_{L \in \mathsf{Zoom}[Q,V]}\left[|N(L) - p_1 m| \geqslant c \cdot p_1 m\right] \leqslant \frac{1}{c^2 p_1 m} \leqslant \frac{q^{r(j-a)}}{c^2 m}.$$

Proof. This is an immediate result of Chebyshev's inequality with the bounds from Claim 5.15. $\square$

Lastly, we use Claim 5.15 to prove Lemma 5.14.

Proof of Lemma 5.14. We have,

$$|\mu(\mathcal{L})-\nu(\mathcal{L})| = \frac{1}{mD}|\sum_{L\in\mathcal{L}} N(L)-p_1m| \leqslant \frac{1}{mD}\sum_{L\in\mathcal{L}}|N(L)-p_1m| \leqslant \frac{|\mathsf{Zoom}[Q,V]|}{mD}\mathop{\mathbb{E}}_{L}[|N(L)-p_1m|],$$

and by Cauchy-Schwartz we get that

$$|\mu(\mathcal{L})-\nu(\mathcal{L})| \leqslant \frac{|\mathsf{Zoom}[Q,V]|}{mD}\sqrt{\mathrm{var}(N(L))}.$$

Plugging in Claim 5.15 and using $|\mathsf{Zoom}[Q,V]| = D/p_1$ we get that

$$|\mu(\mathcal{L})-\nu(\mathcal{L})| \leqslant \frac{1}{\sqrt{p_1m}} \leqslant \frac{3q^{\frac{r}{2}(j-a)}}{\sqrt{m}}. \qquad \square$$

5.3.3 Bound on Maximal Zoom-outs

For this subsection, we work in the second prover's space, $\mathbb{F}_q^V$, and make the assumption that $|V| \gg \ell$, say $|V| \geqslant 2^{100}q^{\ell}$ to be concrete. We first establish several results in the simplified setting where there is no zoom-in. After that we show how to deduce an analogous result with a zoom-in. Throughout this section, we fix T to be a table that assigns to each $L \in \mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)$ a linear function on L.

Definition 5.17. *Given a table T on $\mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)$ and a subspace $Q \subseteq \mathbb{F}_q^V$, we call a zoom-out, function pair, (W, g_W), where $W \subseteq \mathbb{F}_q^V$ and $f : W \to \mathbb{F}_q$, (C,s)-maximal with respect to T on Q if*

$$\Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^V,2\ell)}[g_W|_L \equiv T[L] \mid Q\subseteq L\subseteq W] \geqslant C,$$

and there does not exist another zoom-out function pair, $(W', g_{W'})$ such that $\mathbb{F}_q^V \supseteq W' \supsetneq W$, $g_{W'} : \mathbb{F}'_W \to \mathbb{F}_q$, $g_{W'}|_W \equiv g|_W$ and

$$\Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^V,2\ell)}[g_{W'}|_L \equiv T[L] \mid Q\subseteq L\subseteq W'] \geqslant sC.$$

In the case that $Q = \{0\}$, we say that (W, g_W) is (C,s)-maximal with respect to T.

In the above statement, C should be thought of as small and s should be thought of as an absolute constant. With this in mind, a zoom-out W and a linear function on it g_W is called maximal if there is no zoom-out W' that strictly contains W, and an extension of g_W to $g_{W'}$, that has the same agreement with T as g_W (up to constant factors). As an immediate consequence of the definition of (C,s)-maximal, we have the following lemma, which roughly states that every zoom-out for which there that has a linear function with good agreement inside the zoom-out is contained in a maximal zoom-out (with only slightly worse agreement).

Lemma 5.18. *Let T be a table on $\mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)$, $Q \subseteq \mathbb{F}_q^V$, and $W \subseteq \mathbb{F}_q^V$ be a subspace of codimension r containing Q. Suppose that there exists a linear function $g_W : W \to \mathbb{F}_q$ such that*

$$\Pr_{L \in \mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)}[g_W|_L \equiv T[L] \mid Q \subseteq L \subseteq W] \geqslant C.$$

Then there exists a subspace $W' \supseteq W$ and a linear function $g_{W'} : W' \to \mathbb{F}_q$ such that $g_{W'}|_W \equiv g_W$ and $(g_{W'}, W')$ is (Cs^{-r}, s)-maximal and a linear function $g_{W'} : W' \to \mathbb{F}_q$ such that

$$\Pr_{L \in \mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)}[g_{W'}|_L \equiv T[L] \mid Q \subseteq L \subseteq W] \geqslant C.$$

Proof. This is an immediate consequence of Definition 5.17. If (W, g_W) is (C, s)-maximal then we are done. Otherwise, there must exist W_1, g_{W_1} such that $W_1 \supsetneq W$, $g_{W_1}|_W \equiv g_W$ and

$$\Pr_{L \in \mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)}[g_{W_1}|_L \equiv T[L] \mid Q \subseteq L \subseteq W_1] \geqslant sC.$$

We can repeat this argument at most r times before obtaining some $(g_{W'}, W')$ that is (Cs^{-r}, s) maximal and satisfies $W' \supseteq W$ and $g_{W'}|_W \equiv g_W$. □

The following result, which is key to our analysis gives an upper bound on the number of maximal zoom-outs. This lemma is in fact equivalent to Lemma 1.7, and its proof is deferred to Section 8. In order to present the lemma cleanly, we set the following parameters for the remainder of the section, which can all be considered constant

$$\xi = \delta^5, \quad \delta_2 = \frac{\xi}{100}, \quad t = \left(2^{2+10/\delta_2}\right)! \,.$$

Lemma 5.19. *Let T be a table on $\mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)$ with $|V| \geqslant 2^{100} q^{\ell}$ and set $r \leqslant \frac{10}{\delta}$, $C \geqslant q^{-2(1-\delta^5)\ell}$, and*

$$N \geqslant q^{100(t-1)! r^2 \ell \xi^{-1}}.$$

Suppose that $(W_1, f_1), \dots, (W_N, f_N)$ are zoom-out, function pairs such that the W_i's are all distinct and for each $1 \leqslant i \leqslant m$, $\operatorname{codim}(W_i) = r$, and

$$\Pr_{L \in \mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)}[f_i|_L \equiv T[L] \mid L \subseteq W_i] \geqslant C.$$

Then there is a a subspace V', a linear function $h' : V' \to \mathbb{F}_q$, and a set of subspaces $\mathcal{W}' \subseteq \{W_1, \dots, W_N\}$ of size $N \geqslant q^{50 r \ell \xi^{-1}}$

- *Each $W_i \in \mathcal{W}'$ is strictly contained in V' and has codimension $r' < r$ with respect to V'.*
- *$\mathcal{W}'$ is 2-generic with respect to V'.*
- *For any $W_i \in \mathcal{W}'$, $h'|_{W_i} \equiv f_i$.*

Proof. The proof is deferred to Section 8. □

To bound the number of maximal zoom-outs, we will also need the following list decoding property.

Lemma 5.20. *Let T be a table on $\mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)$, let Q be an r_1-dimensional subspace, and let $W \supseteq Q$ be a subspace of codimension r_2. Suppose that 2ℓ is sufficiently large and $\dim(W) \geqslant 20\ell$. Let $f_1, \ldots, f_m$ be a list of distinct linear functions such that $f_i|_L \equiv T[L]$ for at least β-fraction of the 2ℓ-dimensional subspaces L such that $Q \subseteq L \subseteq W$, for $\beta \geqslant 2q^{-2\ell+r_1} + c$, and $c > 0$. Then,*

$$m \leqslant \frac{4}{c^2} \leqslant \frac{4}{\beta^2}.$$

Proof. The proof is deferred to Appendix D $\square$

Combining Lemma 5.19 with Lemma 5.20, yields a bound on the number of (C, s) maximal zoom-out function pairs with respect to a table T on Q.

Theorem 5.21. *For any table T on $\mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)$ such that $|V| \geqslant q^\ell$, any subspace $Q \subseteq \mathbb{F}_q^V$ of dimension $r_1 \leqslant \frac{10}{\delta}$ and any $C \geqslant q^{-2(1-\delta^3)\ell}$, the number of $(C, \frac{1}{5})$-maximal zoom-out, function pairs with respect to T on Q is at most $\frac{40}{\delta} \cdot C^{-2} \cdot q^{100(t-1)!r^2\ell\xi^{-1}}$.*

Proof. Suppose for the sake of contradiction that $(W_1, f_1), \ldots, (W_M, f_M)$ are $M > \frac{40}{\delta} C^{-2} q^{100(t-1)!r^2\ell\xi^{-1}}$ distinct pairs that are $(C, \frac{1}{5})$ maximal with respect to T on Q. By Lemma 5.20, for each W_i, there are at most $4C^{-2}$ functions $f : W \to \mathbb{F}_q$ such that $f|_L \equiv T[L]$ for at least C-fraction of the $L \in \mathsf{Zoom}[Q, W]$. Thus, there are $C^2 M/4$ distinct W_i's appearing in the pairs, and there is a codimension $r_2 \leqslant \frac{10}{\delta}$ such there are $\frac{M}{\frac{40}{\delta} C^{-2}} = N \geqslant q^{100(t-1)!r^2\ell\xi^{-1}}$ pairs, say, $(W_1, f_1), \ldots, (W_N, f_N)$ zoom-out function pairs that are $(C, \frac{1}{5})$ maximal with respect to T on Q such that the W_i's are all distinct and of codimension r_2 in $\mathbb{F}_q^V$.

Write $\mathbb{F}_q^V = Q \oplus A$. For each $L \subseteq \mathbb{F}_q^V$ of dimension 2ℓ containing Q, there is a unique $L' \subseteq A$ such that $L = Q \oplus L'$. Define the table T' that assigns linear functions to each $L' \in \mathsf{Grass}_q(A, 2\ell - \dim(Q))$ by

$$T'[L'] \equiv T[L' \oplus Q]|_{L'} \tag{8}$$

For each $1 \leqslant i \leqslant N$, let $W_i' \subseteq A$ be the unique subspace such that $W_i = W_i' \oplus Q$. We have that $f_i|_{L'} = T'[L']$ for at least C-fraction of $L' \in \mathsf{Grass}_q(W_i', 2\ell - \dim(Q))$.

By Lemma 5.19 there exists a subspace V', a linear function $h' : V' \to \mathbb{F}_q$, and a subcollection $\mathcal{W}' \subseteq \{W_1', \ldots, W_m'\}$ of size at least $m \geqslant q^{50r\ell\xi^{-1}}$ such that

- Each $W_i' \in \mathcal{W}'$ is has codimension $1 \leqslant r' \leqslant r_2$ with respect to V'.
- $\mathcal{W}'$ is 2-generic with respect to V'.
- For any $W_i' \in \mathcal{W}$, $h'|_{W_i'} = f_i$.

Let $\mathcal{W} = \{W_i \mid W_i' \in \mathcal{W}'\}$ and extend h' to the function $h^\star$ on $V^\star = V' \oplus Q$ so that, $h^\star|_{W_i} \equiv f_i$ for at least q^{-r_1} of the W_i in $\mathcal{W}$. It follows that there is a set $\mathcal{V} = \{W_i \in \mathcal{W} \mid h^\star|_{W_i} \equiv f_i\}$ of size $|\mathcal{V}| \geqslant mq^{-r_2} \geqslant mq^{-\frac{10}{\delta}}$.

Furthermore, because $\mathcal{W}'$ is 2-generic inside of V', $\mathcal{W}$ is 2-generic inside of $V^\star$. We will now finish the proof by applying Lemma 5.14 on $\mathcal{V}$. Specifically, let ν denote the measure over $\mathsf{Zoom}[Q, W^\star]$ generated by choosing $W_i \in \mathcal{V}$ and then $L \in \mathsf{Zoom}[Q, W^\star]$ conditioned on $L \subseteq W_i$, let μ denote the uniform measure over $\mathsf{Zoom}[Q, W^\star]$, and let

$$\mathcal{L} = \{L \subseteq \mathsf{Zoom}[Q, V^\star] \mid h^\star|_L \equiv T[L]\}.$$

Since $h^\star|_{W_i} \equiv f_i$ for every $W_i \in \mathcal{V}$, we have

$$\nu(\mathcal{L}) \geqslant \mathop{\mathbb{E}}_{W_i \in \mathcal{V}}\left[\Pr_{L \in \mathsf{Zoom}[Q,W_i]}[f_i|_L \equiv T[L]]\right] \geqslant C.$$

By Lemma 5.14 with $a = r_1, r = r_2$, and $j = 2\ell$, it follows that

$$\mu(\mathcal{L}) \geqslant \frac{1}{2}\cdot\left(C - \frac{3q^{r_2}(2\ell - r_1)}{\sqrt{m}q^{-r_1}}\right) \geqslant \frac{C}{5}.$$

Summing everything up, this shows that there is a zoom-out function pair $(V^\star, h^\star)$ such that $V^\star \supsetneq W_i$ and $h^\star|_{W_i} \equiv f_i$ for at least one i, and $\Pr_{L\subseteq \mathsf{Grass}_q(\mathbb{F}_q^V,2\ell)}[h^\star|_L \equiv T[L] \mid Q \subseteq L \subseteq V^\star] \geqslant \frac{C}{5}$. This contradicts the assumption that (W_i, f_i) is $(C, \frac{1}{5})$-maximal with respect to T on Q. $\square$

5.4 An Auxiliary Lemma

We conclude this section with an auxiliary lemma that will be used in the analysis.

Lemma 5.22. *Let U be a fixed question to the first prover in the Outer PCP consisting of $3k$-variables in some set of k equations. Let $V \subseteq U$ be a random question to the second prover chosen according to the Outer PCP. Let $W \subseteq \mathbb{F}_q^U$ be a subspace of co-dimension s. Then, with probability at least $1 - 2s\beta$ over the choice of the question V, we have*

$$\dim(W \cap \mathbb{F}_q^V) = |V| - s.$$

Proof. Say that W is given by the constraints $\langle h_1, x\rangle = 0, \ldots, \langle h_s, x\rangle = 0$ for $h_1, \ldots, h_s$ linearly independent. We can view $\mathbb{F}_q^V \subseteq \mathbb{F}_q^U$ as being defined by the constraints $\langle v_{i_1}, x\rangle = 0, \ldots, \langle v_{i_t}, x\rangle = 0$, where $i_1, \ldots, i_t$ correspond to the variables $x_{i_1}, \ldots, x_{i_t}$ in $U \setminus V$. The event $\dim(W \cap \mathbb{F}_q^V) = |V| - s$ is equivalent to the event that $h_1, \ldots, h_b, v_{i_1}, \ldots, v_{i_t}$ are linearly independent. Since $h_1, \ldots, h_s$ are linearly independent, there are s-coordinates such that the restrictions of $h_1, \ldots, h_s$ to these s-coordinates are linearly independent. If none of $i_1, \ldots, i_t$ are in this set of s coordinates, then $h_1, \ldots, h_b, x_{i_1}, \ldots, x_{i_t}$ are linearly independent. Conditioned on the size of $U \setminus V$ being t, this event happens with probability at least $1 - ts/k$, and as the expectation of t is $2\beta k$ it follows that the probability in question is at least $1 - 2s\beta$. $\square$

6 Analysis of the PCP

In this section we show completeness and soundness analysis of the composed PCP construction Ψ from Section 4. As usual, the completeness analysis is straightforward and the soundness analysis will consist the bulk of our effort.

6.1 Completeness

Suppose that the $\mathsf{3Lin}$ instance (X, Eq) has an assignment $\sigma : X \to \mathbb{F}_q$ that satisfies at least $1 - \varepsilon_1$ of the equations in Eq. Let $\mathcal{U}_{\mathsf{sat}} \subseteq \mathcal{U}$ be the set of all $U = (e_1, \ldots, e_k)$ where all k equations $e_1, \ldots, e_k$ are satisfied. Then, $|\mathcal{U}_{\mathsf{sat}}| \geqslant (1 - k\varepsilon_1)\mathcal{U}$. We identify σ with the linear function from $\mathbb{F}_q^X \to \mathbb{F}_q$, assigning the value $\sigma(i)$ to the ith elementary basis element e_i. Abusing notation, we denote this linear map by σ as well.

For each $U \in \mathcal{U}_{\mathsf{sat}}$ and vertex $L \oplus H_U$, we set $T_1[L \oplus H_U] \equiv \sigma|_{L\oplus H_U}$. Since $U \in \mathcal{U}_{\mathsf{sat}}$, these assignments satisfy the side conditions. For all other U's, set $T_1[L \oplus H_U]$ so that the side conditions of H_U are satisfied and $T_1[L \oplus H_U]|_L \equiv \sigma|_L$. Such an assignment is possible because $L \cap H_U = \{0\}$. Similarly, the table T_2 is defined as $T_2[R] \equiv \sigma|_R$.

Sampling a constraint, note that the constraint is satisfied whenever the $L' \oplus H_{U'}$ chosen in step 3 of the test satisfies that $U' \in \mathcal{U}_{\mathsf{sat}}$. As the marginal distribution of $L' \oplus H_{U'}$ is uniform, [8] the distribution of U' is uniform. It follows that the constraint is satisfied whenever $U' \in \mathcal{U}_{\mathsf{sat}}$, which happens with probability at least $1 - k\varepsilon_1$. Thus, $\mathsf{val}(\Psi) \geqslant 1 - k\varepsilon_1$.

6.2 Soundness

In this section we relate the soundness of the composed PCP to that of the outer PCP and prove Lemma 6.1. More precisely, we show:

Lemma 6.1. *For all $\delta > 0$ there are $r \in \mathbb{N}$ and $c(\delta) > 0$ such that the following holds. Let $G^{\otimes k}_{\beta,r}$ parallel repetition of the Smooth Variable versus Equation Game with advice described in Section 3.1.4, and let Ψ be composed PCP described in Section 4. If $\mathsf{val}(G^{\otimes k}_{\beta,r}) < q^{-c(\delta)\cdot O(\ell^2)}$, then $\mathsf{val}(\Psi) \leqslant q^{-2(1-1000\delta)\ell}$.*

The rest of Section 6 is devoted to the proof of Lemma 6.1. The proof heavily relies on the tools from Section 5. Assume, as in lemma Lemma 6.1, that the $\mathsf{val}(G^{\otimes k}_{\beta,r}) < q^{-c(\delta)\cdot O(\ell)^2}$, and suppose for the sake of contradiction that there are tables T_1 and T_2 that are ε-consistent for $\varepsilon \geqslant q^{-2\ell(1-1000\delta)}$. To arrive at a contradiction, we show that this implies strategies for the two provers that with success probability greater than $\mathsf{val}(G^{\otimes k}_{\beta,r})$.

6.2.1 Clique Consistency

To start, we will reduce to the case where T_1 satisfies a condition called *clique-consistency*.

Definition 6.2. *We say an assignment T to $\mathcal{A}$ is clique consistent if for every vertex $L_1 \oplus H_{U_1}$ and for every $L_2 \oplus H_{U_2}, L_3 \oplus H_{U_3} \in [L_1 \oplus H_{U_1}]$, the assignments $T[L_2 \oplus H_{U_2}]$ and $T[L_3 \oplus H_{U_3}]$ satisfy the 1-to-1 constraint between $L_2 \oplus H_{U_2}$ and $L_3 \oplus H_{U_3}$ as specified in Lemma 4.2.*

The following lemma shows that if T_1 and T_2 are ε-consistent assignments to Ψ, then there are clique-consistent assignments $\tilde{T}_1$ and $\tilde{T}_2$ that are also ε-consistent.

Lemma 6.3. *Suppose that the assignments T_1 and T_2 are ε-consistent, then there is a clique-consistent assignment $\tilde{T}_1$ such that $\tilde{T}_1$ and T_2 are ε-consistent.*

Proof. Partition $\mathcal{A}$ into cliques, $\mathcal{A} = \mathsf{Clique}_1 \sqcup \cdots \sqcup \mathsf{Clique}_m$. For each i, choose a random $L \oplus H_U \in \mathsf{Clique}_i$ uniformly, and for every $L' \oplus H_{U'} \in \mathsf{Clique}_i$ assign $\tilde{T}_1[L' \oplus H_{U'}]$ in the unique way that is consistent with $T_1[L \oplus H_U]$ and the side conditions of U' as described in Lemma 4.2. It is clear that $\tilde{T}_1$ is clique consistent, and we next analyze the expected fraction of constraints that $\tilde{T}_1$ and T_2 satisfy.

Note that an alternative description of sampling a constraint in Ψ proceeds as follows. First choose a clique Clique_i with probability that is proportional to by its size, and then choose $L \oplus H_U \in \mathsf{Clique}_i$ in the first step. The rest of the sampling procedure is the same. Let $P(L \oplus H_U)$ be the probability that the test passes conditioned on $L \oplus H_U$ being chosen in the second step. It is clear that every vertex in the clique has equal probability of being chosen, therefore the probability of passing if Clique_i chosen is

$$\frac{1}{|\mathsf{Clique}_i|} \sum_{L \oplus H_U \in \mathsf{Clique}_i} P(L \oplus H_U).$$

[8] This is true because first a clique is chosen with probability that is proportional to its size and then a vertex is sampled uniformly from the clique.

On the other hand, the expected fraction of constraints satisfied by $\tilde{T}_1$ and T_2 (over the randomness of choosing $\tilde{T}_1$) is

$$\sum_{L\oplus H_U} \frac{1}{|\mathsf{Clique}_i|}\cdot P(L\oplus H_U) = \frac{1}{|\mathsf{Clique}_i|}\sum_{L\oplus H_U\in\mathsf{Clique}_i} P(L\oplus H_U).$$

To see this, note that for any $L\oplus H_U$, $\frac{1}{|\mathsf{Clique}_i|}$ is the probability that $T_1[L\oplus H_U]$ is used to define $\tilde{T}_1$ on Clique_i. If this is the case, then the probability the test passes on $\tilde{T}_1$ within Clique_i is $P(L\oplus H_U)$.

Since this holds over every clique, it follows that the expected fraction of constraints satisfied by $\tilde{T}_1$ equals the fraction of constraints satisfied by T_1 and T_2. In particular, there is a choice of $\tilde{T}_1$ such that together with T_2 it satisfies at least ε fraction of the constraints. □

Applying Lemma 6.3 we conclude that there are clique-consistent assignments to Ψ that are ε-consistent, and henceforth we assume that T_1 are clique-consistent to begin with. We remark that, in the notation of Section 4.1.4, the benefit of having a clique-consistent assignment is that the constraint that the verifier checks is equivalent to checking that $T_1[L\oplus H_U]|_R \equiv T_2[R]$. The latter check is a test which in performed within the space $\mathbb{F}_q^U$ of the first prover. We will use this fact in the next section.

6.2.2 A Strategy for the First Prover

Let $p(U)$ be the consistency of T_1 and T_2 conditioned on U being the question to the first prover. As we are assuming that the overall success probability is at least ε, $E_U[p(U)] \geqslant \varepsilon$. By an averaging argument, $p(U)\geqslant \frac{\varepsilon}{2}$ for at least $\frac{\varepsilon}{2}$-fraction of the U's. Call such U's good and let $\mathcal{U}_{\mathsf{good}}$ be the set of good U's.

Let $U\in\mathcal{U}$ be the question to the first prover and let Q be the advice. If $U\notin\mathcal{U}_{\mathsf{good}}$, then the first prover gives up, so henceforth assume that $U\in\mathcal{U}_{\mathsf{good}}$. For such U, the test of the inner PCP passes with probability at least $\frac{\varepsilon}{2}$. More concretely, we have

$$\Pr_{\substack{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell), L\cap H_U=\{0\}\\ \dim(R)=2(1-\delta)\ell}}[T_1[L]|_R \equiv T_2[R] \mid R\subseteq L] \geqslant \frac{\varepsilon}{2}.$$

Next, the first prover chooses an integer $0\leqslant r\leqslant \frac{10}{\delta}$ uniformly, and takes Q to be the span of the first r-advice vectors. By Theorem 5.3, there are r_1, r_2 satisfying $r_1+r_2\leqslant\frac{10}{\delta}$ such that for at least $q^{-6\ell^2}$ of the $Q\subseteq\mathbb{F}_q^U$ of dimension r_1, there exists $W_Q\subseteq\mathbb{F}_q^U$ containing $Q\oplus H_U$ of codimension $r_2\leqslant\frac{10}{\delta}$ and a linear function $g_{Q,W_Q}: W_Q\to\mathbb{F}_q$ that satisfies

$$\Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell), L\cap H_U=\{0\}}[g_{Q,W_Q}|_{L\oplus H_U} \equiv T_1[L\oplus H_U] \mid Q\subseteq L\subseteq W_Q] \geqslant \frac{q^{-2(1-1000\delta^2)\ell}}{5}. \tag{9}$$

For simplicity, set $C = \frac{q^{-2(1-1000\delta^2)\ell}}{5}$. With probability at least $\delta/10$, they choose $r=r_1$, where r_1 is the parameter from Theorem 5.3. Call these dimension r_1 subspaces Q lucky and let $\mathcal{Q}_{\mathsf{lucky}}$ be the set of all lucky $Q\subseteq\mathbb{F}_q^U$. For our analysis, we only analyze the case where the first prover chooses chooses $r=r_1$, which occurs with probability at least $\frac{\delta}{20}$

For each Q such that $Q\in\mathcal{Q}_{\mathsf{lucky}}$ and $Q\cap H_U=\{0\}$, the first prover chooses a W_Q of codimension at most $\frac{10}{\delta}$ and linear function $g_{Q,W_Q}: W_Q\to\mathbb{F}_q$ that satisfies the side conditions on H_U and Equation (9). For such Q that are in $\mathcal{Q}_{\mathsf{lucky}}$ and satisfy $Q\cap H_U=\{0\}$, define

$$\mathcal{L}_Q = \{L\in\mathsf{Zoom}[Q,W_Q] \mid g_{Q,W_Q}|_L \equiv T_1[L\oplus H_U]|_L\}.$$

For Q such that $Q \notin \mathcal{Q}_{\mathsf{lucky}}$ or $Q \cap H_U \neq \{0\}$ define $\mathcal{L}_Q = \emptyset$. Finally, define

$$\mathcal{L} = \{(x_1, \dots, x_{2\ell}) \in \left(\mathbb{F}_q^U\right)^{2\ell} \mid \operatorname{span}(x_1, \dots, x_{r_1}) \in \mathcal{Q}_{\mathsf{lucky}}, \operatorname{span}(x_1, \dots, x_{2\ell}) \in \mathcal{L}_{\operatorname{span}(x_1,\dots,x_{r_1})}\},$$

and let $\mathcal{Q}_{\mathsf{smooth}}$ denote set of r_1-dimensional subspaces Q such that

$$\mathcal{D}'_Q(\mathcal{L}) \geqslant 0.8\mathcal{D}_Q(\mathcal{L}) - \eta^{20}, \tag{10}$$

where $\eta = q^{-\ell^{100}}$. By definition of $\mathcal{D}_Q$ and $\mathcal{L}$, if $Q \in \mathcal{Q}_{\mathsf{lucky}}$, we have

$$\begin{aligned}
\mathcal{D}_Q(\mathcal{L}) &= \Pr_{x=(x_1,\dots,x_{2\ell})\in\mathbb{F}_q^{U\times 2\ell}}[g_{Q,W_Q}|_{\operatorname{span}(x)} \equiv T_1[\operatorname{span}(x) \oplus H_U]|_{\operatorname{span}(x)} \mid \operatorname{span}_{r_1}(x) = Q] \\
&= \Pr_{x=(x_1,\dots,x_{2\ell})\in\mathbb{F}_q^{U\times 2\ell}}[g_{Q,W_Q}|_{\operatorname{span}(x)} \equiv T_1[\operatorname{span}(x \oplus H_U)|_{\operatorname{span}(x)}] \wedge \operatorname{span}(x) \subseteq W_Q \mid \operatorname{span}_{r_1}(x) = Q] \\
&= \Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell), L\cap H_U=\{0\}}[g_{Q,W_Q}|_{L\oplus H_U} \equiv T_1[L \oplus H_U] \wedge L \subseteq W_Q \mid Q \subseteq L]
\end{aligned}$$

where the second transition is because, by definition, every $L \in \mathcal{L}_Q$ is contained in W_Q. For the third transition, we are ignoring the probability that the x_i's are linearly dependent and $\operatorname{span}(x) \cap H_U \neq \{0\}$. Indeed, the probability that either of these events occur is at most $q^{2\ell-3k} + q^{2\ell-2k}$, and is negligible anyways.

Continuing, for $Q \in \mathcal{Q}_{\mathsf{lucky}}$, we have

$$\begin{aligned}
\mathcal{D}_Q(\mathcal{L}) &= \Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell), L\cap H_U=\{0\}}[g_{Q,W_Q}|_{L\oplus H_U} \equiv T_1[L \oplus H_U] \wedge L \subseteq W_Q \mid Q \subseteq L] \\
&= \Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell), L\cap H_U=\{0\}}[g_{Q,W_Q}|_{L\oplus H_U} \equiv T_1[L \oplus H_U] \mid Q \subseteq L \subseteq W_Q] \\
&\cdot \Pr_{L\in\mathsf{Grass}_q(\mathbb{F}_q^U,2\ell), L\cap H_U=\{0\}}[L \subseteq W_Q \mid Q \subseteq L] \\
&\geqslant q^{-r_2(2\ell-r_1)} \cdot C,
\end{aligned} \tag{11}$$

where in the third transition uses Equation (9) to lower bound the first term by C.

If either $Q \notin \mathcal{Q}_{\mathsf{lucky}}$, $Q \cap H_U \neq \{0\}$, or $Q \notin \mathcal{Q}_{\mathsf{smooth}}$, then the first prover gives up. Otherwise the prover extends the function g_Q to a linear function on the entire space $\mathbb{F}_q^U$ randomly, and we denote this extension by $g : \mathbb{F}_q^U \to \mathbb{F}_q$. The prover outputs the string $s_{Q,U}$ as their answer where $s_{Q,U} \in \mathbb{F}_q^U$ is the unique string such that $g(x) = \langle s_{Q,U}, x\rangle$ for all $x \in \mathbb{F}_q^U$. As g_{Q,W_Q}, and by extension g, respects the side conditions, it follows that $s_{Q,U}$ satisfies the k-linear equations of U.

6.2.3 A Strategy for the Second Prover

Let V be the question to the second prover. The second prover will use a table $\tilde{T}_1$ to derive their strategy. The table $\tilde{T}_1$ is obtained from T_1 as follows. For a question V to the second prover, let $U \supseteq V$ be an arbitrary question to the first prover. For all 2ℓ-dimensional subspaces $L \subseteq \mathbb{F}_q^V$, define

$$\tilde{T}_1[L] \equiv T_1[L \oplus H_U]|_L.$$

In order to make sure that $\tilde{T}_1$ is well defined, we note two things. First, the subspace $L \oplus H_U$ can be viewed as a subspace of $\mathbb{F}_q^U$ because each $L \subseteq \mathbb{F}_q^V$ can be "lifted" to a subspace of $\mathbb{F}_q^U$ by inserting 0's into the coordinates corresponding to $V \setminus U$. Second, note that the choice of $U \supseteq V$ does not matter when defining

$\tilde{T}_1[L]$. Indeed, for a fixed L, the vertices $L \oplus H_U$ over all $U \supseteq V$ are in the same clique. Since T_1 is clique consistent, it does not matter which U is chosen when defining $\tilde{T}_1[L]$, as all choices lead to the same function $T_1[L \oplus H_U]|_L$. Therefore the second prover can construct the table $\tilde{T}_1$.

After constructing $\tilde{T}_1$, the second prover then chooses a dimension $0 \leqslant r \leqslant \frac{10}{\delta}$ uniformly for the advice Q. Note that with probability at least $\frac{\delta}{20}$ the second prover also chooses $r = r_1$. The second prover then chooses a zoom-out function pair $(W_{\mathsf{second}}, g_{Q,W_{\mathsf{second}}})$ that is

$$\left(\frac{C}{4 \cdot 5^{r_2}}, \frac{1}{5}\right)$$

-maximal with respect to $\tilde{T}_1$ on Q if one exists (and gives up otherwise).

Finally, the second prover extends the function $g_{Q,W_{\mathsf{second}}}$ randomly to a linear function on $\mathbb{F}_q^V$ to arrive at their answer. The resulting function is linear and it is equal to the inner product function $y \to \langle s_{Q,V}, y\rangle$ for some unique string $s_{Q,V} \in \mathbb{F}_q^V$. The second prover outputs $s_{Q,V}$ as their answer.

6.2.4 The Success Probability of the Provers

In order to be successful, a series of events must occur. We go through each one and state the probability that each occurs. At the end this yields a lower bound on the provers' success probability. We remark that the analysis of this sections requires Lemmas 5.4 and 5.6, so recall that k and β are set according to Equation (4) in Section 1.2.6 so that these lemmas hold.

First, the provers need $U \in \mathcal{U}_{\mathsf{good}}$, which occurs with probability at least $\frac{\varepsilon}{2}$. Assuming that this occurs, the provers then both need to choose $r = r_1$ for the dimension of their zoom-in, which happens with probability at least $\frac{\delta^2}{101}$. If both provers choose $r = r_1$, they both receive advice Q as the span of r_1 random vectors.

The provers then need $Q \in \mathcal{Q}_{\mathsf{lucky}}$, $Q \cap H_U = \{0\}$, and $Q \in \mathcal{Q}_{\mathsf{smooth}}$. When analyzing the probability that these three events occur, we need to recall that the advice vectors are actually drawn according to distribution $\mathcal{D}'_{r_1}$, the distribution described in Section 5.2.3. We will analyze the probability that the three events occur under $\mathcal{D}_{r_1}$ and then appeal to the covering property of Lemma 5.6. By Theorem 5.2, the first item occurs with probability at least $q^{-6\ell^2}$. On the other hand the probability that the second item does not occur is at most $\sum_{i=0}^{r_1} \frac{q^i q^k}{q^{3k}} \leqslant q^{r_1+1-2k}$, while the probability that the third item does not occur is at most η^{20} by Lemma 5.5. Altogether we get that with probability at least

$$q^{-6\ell^2} - q^{r^*+1-2k} - \eta^{20} q^{-7\ell^2}$$

under $\mathcal{D}_{r_1}$, we have $Q \in \mathcal{Q}_{\mathsf{lucky}}$, $Q \cap H_U = \{0\}$, and $Q \in \mathcal{Q}_{\mathsf{smooth}}$. By Lemma 5.6, we have that $Q \in \mathcal{Q}_{\mathsf{lucky}}$, $Q \cap H_U = \{0\}$, and $Q \in \mathcal{Q}_{\mathsf{smooth}}$ with probability at least $q^{-8\ell^2}$ under $\mathcal{D}'_{r_1}$ - the distribution which the Q is actually drawn from.

Now let us assume that $U \in \mathcal{U}_{\mathsf{good}}$, and both provers receive an r_1-dimensional advice Q such that $Q \in \mathcal{Q}_{\mathsf{lucky}}$, $Q \cap H_U = \{0\}$, and $Q \in \mathcal{Q}_{\mathsf{smooth}}$. The first prover chooses the function $g_Q : W_Q \to \mathbb{F}_q$. Write $\mathrm{codim}(W_Q) = r_2$. Since $Q \in \mathcal{Q}_{\mathsf{lucky}}$, by Equation (11) we have

$$\mathcal{D}_Q(\mathcal{L}) \geqslant q^{-r_2(2\ell - r_1)} \cdot C,$$

since by definition $\mathcal{L}$ contains at least C-fraction of $L \in \mathsf{Zoom}[Q, W]$, which is in turn at least $q^{-r_2(2\ell-r_1)}$ fraction of $L \in \mathsf{Zoom}[Q, \mathbb{F}_q^U]$. Next because $Q \in \mathcal{Q}_{\mathsf{smooth}}$, we have

$$\mathcal{D}'_Q(\mathcal{L}) \geqslant 0.8 \cdot \mathcal{D}_Q(\mathcal{L}) - \eta^{60} \geqslant q^{-r_2(2\ell-r_1)} \cdot \frac{C}{2}, \tag{12}$$

by Equation (10).

Now let $W_Q[V] = W_Q \cap \mathbb{F}_q^V$. By Lemma 5.22, with probability at least $1 - 2\beta r_2$ we have that $\mathrm{codim}(W_Q[V]) = r_1$ inside of $\mathbb{F}_q^V$. Combining this with an averaging argument on Equation (12), we have that with probability at least $\frac{C}{5} - 2\beta r_2 \geqslant \frac{C}{6}$ over V,

$$\begin{aligned}\Pr_{x_i \in \mathbb{F}_q^V, w_i \in H_U}[\mathrm{span}(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell}) \in \mathcal{L} \mid \ \mathrm{span}_{r_1}(x) = Q] = \mathcal{D}'_Q(\mathcal{L}) \\ \geqslant q^{-r_2(2\ell - r_1)} \cdot \frac{C}{4},\end{aligned} \tag{13}$$

and $\mathrm{codim}(W_Q[V]) = r_1$. We call such V consistent. In the probability above, and henceforth, we view vectors $x_i \in \mathbb{F}_q^V$ as vectors in $\mathbb{F}_q^U$ with 0's appended in the missing coordinates. At this point there is the slight issue that $\mathrm{span}(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell})$ does not actually correspond to a random entry of the second prover's table, $\tilde{T}_1$. Indeed, the second prover can only choose $(x_1, \dots, x_{2\ell}) \in \mathbb{F}_q^V$, choose some question U to the first prover that contains V, lift these to $\mathbb{F}_q^U$ by inserting zeros in the missing coordinates, and look at the entry $L \oplus H_U$ where $L = \mathrm{span}(x_1, \dots, x_{2\ell})$. They do not know the question U and the side conditions H_U, and hence could not sample the $w_i \in H_U$. However, notice that for any $w_1, \dots, w_{2\ell} \in H_U$, we have that,

$$\mathrm{span}(x_1, \dots, x_{2\ell}) \oplus H_U = \mathrm{span}(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell}) \oplus H_U.$$

We can thus view the $w_1, \dots, w_{2\ell} \in H_U$ being sampled and added to $x_1, \dots, x_{2\ell}$ as a virtual step. In the next two equations, we ignore the probability that $(x_1, \dots, x_{2\ell})$ or $(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell})$ are not linearly independent or have spans intersecting H_U to make the expressions above simpler. This probability is at most $\frac{q^{2\ell} + q^{k+2\ell}}{q^{3k}}$ and is negligible anyways. We have,

$$\begin{aligned}&\Pr_{L \in \mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)}[\tilde{T}_1[L] = g_Q|_L \mid Q \subseteq L \subseteq W_Q[V]] \\ &\qquad = \Pr_{L = \mathrm{span}(x_1, \dots, x_{2\ell})}[\tilde{T}_1[L] = g_Q|_L \mid Q \subseteq L \subseteq W_Q[V]] \\ &\qquad = \Pr_{L = \mathrm{span}(x_1, \dots, x_{2\ell})}[T_1[L \oplus H_U] = g_Q|_L \mid Q \subseteq L \subseteq W_Q[V]] \\ &\qquad = \Pr_{L = \mathrm{span}(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell})}[T_1[L \oplus H_U] = g_Q|_L \mid Q \subseteq L \subseteq W_Q[V]] \\ &\qquad = \Pr_{L = \mathrm{span}(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell})}[\tilde{T}_1[L] = g_Q|_L \mid Q \subseteq L \subseteq W_Q[V]].\end{aligned}$$

This last probability can be related to $\mathcal{D}'_Q(\mathcal{L})$

$$\begin{aligned}&\Pr_{L = \mathrm{span}(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell})}[\tilde{T}_1[L] = g_Q|_L \mid Q \subseteq L \subseteq W_Q[V]] \\ &\geqslant q^{r_2(2\ell - r_1)} \cdot \Pr_{x_i \in \mathbb{F}_q^V, w_i \in H_U}[(x_1 + w_1, \dots, x_{2\ell} + w_{2\ell}) \in \mathcal{L} \mid \ \mathrm{span}_{r_1}((x_i + w_i)) = Q] \\ &= q^{r_2(2\ell - r_1)} \cdot \mathcal{D}'_Q(\mathcal{L}) \\ &\geqslant \frac{C}{4}.\end{aligned}$$

By Lemma 5.18, there exists some $(W'_Q[V], g_{Q, W'_Q[V]})$ that is $\left(\frac{C}{4 \cdot 5^{r_2}}, \frac{1}{5}\right)$-maximal and satisfies $W'_Q[V] \supseteq W_Q[V]$, $g_{Q, W'_Q[V]} : W'_Q[V] \to \mathbb{F}_q$ is linear, and $g_{Q, W'_Q[V]}|_{W[V]} = g_Q|_{W_Q[V]}$. By Theorem 5.21, the number of $\left(\frac{C}{4\cdot 5^{r_2}}, \frac{1}{5}\right)$ is at most

$$M = \frac{40}{\delta} \cdot 5^{2r_2+2} C^{-2} \cdot q^{100(t-1)! r^2 \ell \xi^{-1}} \leqslant q^{c(\delta)\ell},$$

where $c(\delta)$ is some function depending only on δ. Thus, the second prover chooses $(W'_Q[V], g_{Q,W'_Q[V]})$ with probability at least $\frac{1}{M}$. Finally, if the second prover chooses $(W'_Q[V], g_{Q,W'_Q[V]})$, then the provers succeed if both provers extend their functions, $g_Q|_{W[V]}$ and $g_{Q,W'_Q[V]}$ in the same manner. This occurs with probability at least $q^{-\operatorname{codim}(W[V])} \geqslant q^{-10/\delta}$.

Putting everything together, we get that the provers succeed with probability at least

$$\frac{\varepsilon}{2} \cdot \frac{\delta^2}{101} \cdot q^{-8\ell^2} \cdot \frac{C}{5} \cdot \frac{1}{M} \cdot q^{-10/\delta} = q^{-c(\delta)\cdot O(\ell^2)},$$

where the first term is the probability that $U \in \mathcal{U}_{\mathsf{good}}$, the second term is the probability that both provers choose the same zoom-in dimension, the third term is the probability that $Q \in \mathcal{Q}_{\mathsf{lucky}}$, $Q \cap H_U = \{0\}$, $Q \in \mathcal{Q}_{\mathsf{smooth}}$, the fourth term is the probability that V is consistent, the fifth term is the probability that the second prover chooses the a function that extends the first prover's answer, and the final term is the probability that both provers extend their functions in the same manner.

This proves Lemma 6.1.

7 Proofs of the Main Theorems

7.1 Proof of Theorem 1.3

Theorem 1.3 follows by applying out PCP construction from Section 4 starting with an instance of 3-Lin from Theorem 2.1. We may take $q = 2$, fix $\delta > 0$ to be a small constant and then take ℓ sufficiently large compared to δ^{-1}, then k and β according to Equation (4), and finally take the completeness of the 3-Lin instance, $1 - \eta$, so that $\eta < 1/k$. It follows that if the original 3-Lin instance is at least $1 - \eta$ satisfiable, then $\mathsf{val}(\Psi) \geqslant 1 - k\eta$. On the other hand, if the if the original instance is at most s satisfiable for some constant $s > 0$, then by Claim 3.2, the value of the outer PCP is at most

$$\mathsf{val}(G_{\beta,r}^{\otimes k}) \leqslant 2^{-\Omega\left((1-s)^2 q^{-r+\frac{2c}{3}\ell}\right)} < q^{-c(\delta)O(\ell)^2},$$

since we take ℓ sufficiently large compared to δ^{-1}. By Lemma 6.1, it follows that if the original instance is at most s satisfiable, then $\mathsf{val}(\Psi) \leqslant q^{-2(1-1000\delta)\ell}$. The proof is concluded as the alphabet size of Ψ is $O(q^{2\ell})$.

7.2 Proof of Theorem 1.4

To show quasi-NP-hardness for approximate Quadratic Programming, we rely on the following result due to [ABH+05], who show a reduction from 2-Prover-1-Round Games to Quadratic Programming.

Theorem 7.1. *There is a reduction from a 2-Prover-1-Round Games, Ψ with graph $G = (L \cup R, E)$ and alphabets Σ_L, Σ_R to a Quadratic Programming instance A such that:*

- *The running time of the reduction and the number of variables in A is polynomial in $|L| + |R|$ and $2^{|\Sigma_L|}$.*

- *If* $\mathsf{val}(\Psi) \geqslant 1-\eta$*, then* $\mathsf{OPT}(A) \geqslant 1-\eta-\frac{1}{|L|+|R|}$.
- *If* $\mathsf{val}(\Psi) \leqslant \varepsilon$*, then* $\mathsf{OPT}(A) \leqslant O(\varepsilon)$.

We are now ready to prove Theorem 1.4.

Proof of Theorem 1.4. Starting with a SAT instance of size n, which may be arbitrarily large, we take the instance of $\mathsf{Gap3Lin}$ from Theorem 2.2 of size $N \leqslant 2^{O(\log^2 n)}$ and field size $q = \Theta(\log n)$ which is a power of 2 as the starting point of our reduction. Take $\delta > 0$ to be a small constant, and ℓ to be a sufficiently large constant relative to δ^{-1} in our composed PCP. Finally, we pick k and β by Equation (4). This yields a $2^{O(k \log^2 n)}$-time reduction from SAT to a 2-Prover-1-Game on $G = (L \cup R, E)$, with alphabets Σ_L, Σ_R and the following properties:

- $|R| + |L| = O(N^k \cdot q^{3k+2\ell})$.
- $|\Sigma_R| \leqslant |\Sigma_L| = q^{2\ell}$.
- The completeness is at least $1 - k\eta$, where $\eta = 2^{-\Theta(\sqrt{\log n})}$.
- The soundness is at most $q^{-2(1-1000\delta)\ell}$.

Indeed, the first 3 properties are clear. For the soundness, as the original 3-Lin instance is at most $1-\varepsilon$ satisfiable for $\varepsilon = 1/\log^3 N$, we get from Claim 3.2 that

$$\mathsf{val}(G_{\beta,r}^{\otimes k}) \leqslant 2^{-\Omega(\varepsilon^{-2} q^{-r+\frac{2c\ell}{3}})} \leqslant q^{-c(\delta)O(\ell)^2},$$

as ℓ is sufficiently large relative to δ^{-1}, so the soundness of the composed PCP follows by Lemma 6.1. Applying the reduction of Theorem 7.1, we get a reduction to a Quadratic Programming instance A such that,

- The running time of the reduction and number of variables in A are both
$$M = \mathrm{poly}(2^{O(\log^2 n) q^{2(1+c)\ell}} (\log n)^{O(q^{2(1+c)\ell})} 2^{q^{2\ell}}).$$
- If the original SAT instance is satisfiable, then
$$\mathsf{OPT} \geqslant 1 - 2^{-\Omega(\sqrt{\log n})}.$$
- If the original SAT instance is not satisfiable, then
$$\mathsf{OPT} \leqslant O(q^{-2(1-1000\delta)\ell}).$$

Note that
$$\log(M) = q^{2(1+c)\ell} O(\log^2 n),$$
whereas the gap between the satisfiable and unsatisfiable cases is $\Omega\left(q^{-2(1-1000\delta)\ell}\right) = \frac{1}{\log(M)^{1-O(\delta)}}$. Altogether, this shows that for all $\varepsilon > 0$ there is $C > 0$ such that unless $\mathsf{NP} \subseteq \mathsf{DTIME}\left(2^{\log(n)^C}\right)$, there is no $\log(M)^{1-\varepsilon}$-approximation algorithm for Quadratic Programming on M variables. $\square$

7.3 Proof of Theorem 1.5

In this section we prove Theorem 1.5, and for that we must first establish a version of Theorem 1.3 for bi-regular graphs of bounded degree. The proof of this requites minor modifications of our construction, as well as the right degree reduction technique of Moshkovitz and Raz [MR10].

7.3.1 Obtaining a Hard Instance of Bipartite Biregular 2-CSP

In this section first we show that the 2-Prover-1-Round game from Theorem 1.3 can be transformed into a hard instance of biregular, bipartite 2-CSP with bounded degrees. This version may be useful for future applications, and is formally stated below. Call a bipartite 2-CSP (d_1, d_2)-regular if the left degrees of its underlying graph are all d_1, and the right degrees of its underlying graph are all d_2.

Theorem 7.1. *For every $\varphi, \varepsilon > 0$, and sufficiently large $R \in \mathbb{N}$, there exist $d_1, d_2 \in \mathbb{N}$ such that given a bipartite (d_1, d_2)-regular 2-CSP, Ψ, with alphabet size R, it is NP-hard to distinguish the following two cases:*

- *Completeness:* $\mathsf{val}(\Psi) \geqslant 1 - \varphi$,
- *Soundness:* $\mathsf{val}(\Psi) \leqslant \frac{1}{R^{1-\varepsilon}}$.

To prove Theorem 7.1, we start with an instance Ψ from Theorem 1.3, and first argue that Ψ can be made left regular while almost preserving soundness and completeness by deleting a small fraction of left vertices. We then use the right degree reduction technique of Moshkovitz and Raz [MR10], to obtain a bounded degree bi-regular bipartite 2-CSP.

Fix $\varphi, \varepsilon > 0$, and let Ψ to be the 2-Prover-1-Round game constructed for Theorem 1.3. Recall that this requires us to choose some large enough ℓ relative to $\varphi^{-1}, \varepsilon^{-1}$, some large enough q relative to ℓ, and set $R = q^{2\ell}$. We also set $\delta = \frac{\varepsilon}{1000}$, $0 < c$ arbitrarily small relative to δ, and $k = q^{2(1+c)\ell}$. Finally, we construct our 2-Prover-1-Round game from a hard instance of $\mathsf{Gap3Lin}$ with the appropriate completeness and soundness, so that it is NP-hard to distinguish between,

$$\mathsf{val}(\Psi) \geqslant 1 - \varphi \quad \text{and} \quad \mathsf{val}(\Psi) \leqslant \frac{1}{q^{2(1-1000\delta)\ell}} = \frac{1}{R^{1-\varepsilon}}.$$

It is clear that our 2-Prover-1-Round game can equivalently be viewed as an instance of bipartite 2-CSP, so let us analyze the underlying graph. Let $\mathcal{U}$ denote the set of possible questions to the first prover. Recall that the set of left vertices is,

$$\mathsf{Left} = \{L \oplus H_U \mid \forall U \in \mathcal{U}, \forall L \in \mathsf{Grass}_q(\mathbb{F}_q^U, 2\ell), L \cap H_U = \{0\}\},$$

while the set of right vertices is

$$\mathsf{Right} = \{R \in \mathsf{Grass}_q(\mathbb{F}_q^U, 2(1-\delta)\ell) \mid U \in \mathcal{U}\}.$$

The edges and constraints of this graph are generated by a randomized process. Equivalently, there is a weight function $w(\cdot)$ over edges $(L \oplus H_U, R)$. Recall that the weighting is defined by first choosing a uniform $L \oplus H_U \in \mathsf{Left}$, and then $R \in \mathsf{Right}$ according to the process descrbied in Section 4.1.4. For a fixed $L \oplus H_U$, define

$$w_{L \oplus H_U}(R) = \frac{|\{L' \oplus H_{U'} \in [L \oplus H_U] \mid R \subseteq L'\}|}{|[L \oplus H_U]|} \cdot \frac{1}{\begin{bmatrix} 2\ell \\ 2(1-\delta)\ell \end{bmatrix}_q}.$$

This is the probability of choosing the $(L \oplus H_U, R)$ conditioned on first choosing $L \oplus H_U$. Since we choose $L \oplus H_U \in \mathsf{Left}$ uniformly, it follows that

$$w(L \oplus H_U, R) = \frac{w_{L\oplus H_U}(R)}{|\mathsf{Left}|}.$$

Define the neighborhood of a vertex as,

$$\mathsf{nb}(L \oplus H_U) = \{R \in \mathsf{Right} \mid w_{L\oplus H_U}(R) > 0\}.$$

We will now attempt to remove some left vertices and obtain a bipartite, left-regular 2-CSP. To this end, we call $L \oplus H_U$ trivial if there is an equation $e \in U$ such that for every basis of $x_1, \ldots, x_{2\ell} \in \mathbb{F}_q^{U_i}$ of L, the points x_i restricted to the variables in e are of the form (α, α, α) for some $\alpha \in \mathbb{F}_q$.

Claim 7.2. *The fraction of $L \oplus H_U \in \mathsf{Left}$ that are trivial is at most $2q^{-(2-2c)\ell}$.*

Proof. Fix a $U \in \mathcal{U}$. Note that it suffices to show that at most $2q^{-(2-2c)\ell}$ vertices of the form $L \oplus H_U$ are trivial, as for each $U \in \mathcal{U}$, there are an equal number of vertices $L \oplus H_U$.

Write $U = (x_1, \ldots, x_{3k})$, where the ith equation in U contains the variables $x_{3i-2}, x_{3i-1}, x_{3i}$. Call these three coordinates a block, so that each $x \in \mathbb{F}_q^{3k}$ consists of k blocks of consecutive coordinates. Let us bound the fraction of L such that $L \oplus H_U$ is trivial. For $y_1, \ldots, y_{2\ell} \in \mathbb{F}_q^{3k}$, let $s(y_1, \ldots, y_{2\ell})$ be the number of blocks where $y_1, \ldots, y_{2\ell}$ are all of the form (α, α, α) for some $\alpha \in \mathbb{F}_q$. Then

$$\Pr_L[L \oplus H_U \text{ is trivial}] \leqslant 2 \Pr_{y_1,\ldots,y_{2\ell}}[s(y_1, \ldots, y_{2\ell}) = 0],$$

where the factor of 2 accounts for the probability that either $y_1, \ldots, y_{2\ell}$ are not linearly dependent, or $\mathrm{span}(y_1, \ldots, y_{2\ell}) \cap H_U \neq \{0\}$. Note that the probability that a specific block is trivial is $q^{-4\ell}$, hence by linearity of expectation we get that

$$\mathop{\mathbb{E}}_{y_1,\ldots,y_{2\ell}}[s(y_1, \ldots, y_{2\ell})] = kq^{-4\ell} = q^{-(2-2c)\ell},$$

and therefore

$$\Pr_{x_1,\ldots,x_{2\ell}}[s(x_1, \ldots, x_{2\ell}) \geqslant 1] \leqslant q^{-(2-2c)\ell}.$$

□

Let Ψ' be the instance obtained from Ψ after removing all trivial $L \oplus H_U$ from Ψ, so that the new instance Ψ' does not contain any trivial vertices. Let Left' denote the set of left vertices in Ψ' and let $w'(\cdot)$ denote the weight function over edges in Ψ', which is given by choosing $L \oplus H_U \in \mathsf{Left}'$ uniformly, and then choosing $R \in \mathsf{nb}(L \oplus H_U)$ with probability proportional to $w_{L\oplus H_U}(R)$. It follows that,

$$w'(L \oplus H_U, R) = \frac{w_{L\oplus H_U}(R)}{|\mathsf{Left}'|}.$$

Let E' be a set of edges in Ψ and for notational purposes let us write $w'(L \oplus H_U, R) = 0$ if $L \oplus H_U$ is trivial and not in Ψ'. We have

$$w(E) - 4q^{-2+2c\ell} \leqslant w'(E') \leqslant \frac{w(E')}{1 - 2q^{-2+2c\ell}}. \tag{14}$$

The upper bound is clear from Claim 7.2. For the lower bound, we have $w(E') \geqslant \frac{w(E') - 2q^{-2+2c\ell}}{1-2q^{-2+2c\ell}} \geqslant w(E) - 4q^{-2+2c\ell}$.

It follows that Ψ' has completeness at least $1 - \psi - 4q^{-(2-2c)\ell}$ and soundness at most $\frac{2}{q^{-2(1-1000\delta)\ell}}$. We now bound the size of the neighborhoods in Ψ'.

Claim 7.3. *For each $L \oplus H_U$, we have*

$$|\mathsf{nb}(L \oplus H_U)| \leqslant 10^k q^{6k\ell}.$$

Proof. Let $U = (x_1, \ldots, x_{3k}) = (e_1, \ldots, e_k)$ and suppose equation e_i contains variables $(x_{3i-2}, x_{3i-1}, x_{3i})$. Since L is not trivial, for each i, there must be a point $v \in L \subseteq \mathbb{F}_q^U$ such that the values of v restricted to the coordinates of variables $(x_{3i-2}, x_{3i-1}, x_{3i})$ are not all equal. Without loss of generality, say that it is x_{3i} for each $1 \leqslant i \leqslant k$. It follows that in order to have

$$L \oplus H_U \oplus H_{U'} = L' \oplus H_U \oplus H_{U'},$$

U' must contain an equation with the variable x_{3i} for each $1 \leqslant i \leqslant k$. Let E_i denote this set of equations for each i. By the regularity assumptions on our $\mathsf{3Lin}$ instance, $|E_i| \leqslant 10$ and $E_i \cap E_j = \emptyset$ for $i \neq j$. It follows that U' must contain exactly one equation from each E_i, and that these form all k equations of U', so there are at most 10^k possible U' for which there can exist $L' \subseteq U'$, such that $L' \oplus H_{U'} \in [L \oplus H_U]$. The lemma follows from the observation that $|\mathsf{Grass}_q(3k, 2(1-\delta)\ell)| \leqslant q^{6k\ell}$. □

Performing the same procedure as in [KR03, Lemma 3.4], we can turn Ψ'' into a bipartite, left-regular 2-CSP instance without losing too much in completeness or soundness. Let $Q = 10^k q^{6k\ell}$ be the upper bound on neighborhood sizes in Claim 7.3.

Claim 7.4. *For any $C \in \mathbb{N}$, there is a polynomial time algorithm that takes Ψ' as input and outputs a bipartite 2-CSP Ψ'' that is left regular with degree $C \cdot Q$, and has*

- *Completeness $1 - \varphi - 4q^{-(2-2c)\ell} - \frac{1}{C}$.*
- *Soundness $\frac{2}{q^{-2(1-1000)\delta\ell}} + \frac{1}{C}$.*

Proof. We define Ψ'' as follows. For each vertex $L \oplus H_U \in \mathsf{Left}'$, do the following. Let $R_1, \ldots, R_m$ be the vertices in $\mathsf{nb}(L \oplus H_U)$. For each $2 \leqslant i \leqslant m$, add $\lfloor w_{L\oplus H_U}(R_i) C \cdot Q \rfloor$ edges from $L \oplus H_U$ to R_i. Also add $C \cdot Q - \sum_{i=2}^{m} \lfloor w_{L\oplus H_U}(R_i) C \cdot Q \rfloor$ edges from $L \oplus H_U$ to R_1. It is clear that Ψ'' is left regular with degree $C \cdot Q$, and that for each R_i, there are at most $w_{L\oplus H_U}(R_i) C \cdot Q$ edges between $L \oplus H_U$ and R_i for $2 \leqslant i \leqslant m$, and at most $\left(w_{L\oplus H_U}(R_i) + \frac{1}{C}\right) C \cdot Q$ edges between $L \oplus H_U$ and R_1.

For the completeness and soundness, consider a left vertex $L \oplus H_U$. Then it is clear by the above that if a labelling satisfies $1 - c$ fraction of constraints involving $L \oplus H_U$ in Ψ', then in Ψ'' the same labelling satisfies at least $\left(1 - c - \frac{1}{C}\right) C \cdot Q$ of the edges incident to $L \oplus H_U$. Similarly, if a labelling satisfies at most s fraction of constraints involving $L \oplus H_U$ in Ψ', then it satisfies at most $\left(1 - c + \frac{1}{C}\right) C \cdot Q$ edges involving $L \oplus H_U$. □

Applying Claim 7.4 with $C = q^{10\ell}$, we obtain a bipartite 2-CSP, Ψ'', that is left regular with degree $q^{10\ell} Q$, that still has nearly the same completeness and soundness as our original instance Ψ. We will now create a 2-regular bipartite CSP from Ψ'', by using the degree reduction technique of Moshkovitz and Raz [MR10].

Lemma 7.5. *[MR10] For any parameter d, there is a polynomial time algorithm that takes a bipartite, left regular 2-CSP with left degree, d_{left}, completeness $1 - \varphi'$ and soundness s, and outputs a bipartite, $(d_{\mathsf{left}}, dd_{\mathsf{left}})$-regular 2-CSP completeness $1 - \varphi'$ and soundness $s + O\left(d^{-1/2}\right)$.*

We are now ready to complete the proof of Theorem 7.1.

Proof of Theorem 7.1. Applying, Lemma 7.5 to Ψ'' with $d = q^{10\ell}$, we obtain a bipartite $(q^{10\ell}Q, q^{20\ell}Q)$-regular 2-CSP, Ψ_0, with completeness at least $1 - \varphi - 5q^{-(2-2c)\ell}$ and soundness at most $\frac{3}{q^{-2(1-1000\delta)\ell}}$. By setting q and ℓ large enough relative to φ, Ψ'' has completeness $1 - 1.1\varphi$. The soundness is at most, $\frac{1}{R^{1-0.9\varepsilon}}$. As the original φ and ε can be arbitrarily small positive constants, Theorem 7.1 follows. $\square$

7.3.2 Sparsification

In [LM23], Lee and Manurangsi show how to conclude Theorem 1.5 from Theorem 7.1 via a sparsification procedure. We summarize the steps here. Fix the $\eta > 0$ for Theorem 1.5. Set $\varphi = \varepsilon = 0.01\eta$ in Theorem 7.1 and let Ψ be the resulting hard bipartite (d_1, d_2)-regular 2-CSP, and it is NP-hard to distinguish between.

$$\mathsf{val}(\Psi) \geqslant 1 - \varphi \quad \text{and} \quad \mathsf{val}(\Psi) \leqslant \frac{1}{R^{1-\varepsilon}}.$$

Now, observe that the degrees of Ψ can be multiplied by arbitrary constants by copying vertices.

Lemma 7.6. *[LM23, Lemma 10] For any integers d_1, d_2, c_1, c_2, there is a polynomial time reduction from a bipartite (d_1, d_2)-biregular CSP, Ψ, to a bipartite $(c_2d_1d_2, c_1d_1d_2)$-biregular CSP Ψ', such that $\mathsf{val}(\Psi) = \mathsf{val}(\Psi')$, and such that the left and right alphabet sizes are preserved.*

It is then shown in [LM23] that one can perform a subsampling procedure to Ψ', that significantly lowers the degree, while not increasing the soundness or alphabet size too much.

Theorem 7.7. *[LM23, Theorem 11] For any $0 < \nu_1 < \nu_2 \leqslant 1$, such that any positive integer C, and any sufficiently large positive integers $d_A, d_B \geqslant d_0(\varphi, \nu)$, and $R \geqslant R_0(\delta, \nu, d_A, d_B)$, the following holds: there is a randomized polynomial-time reduction from a bipartite (d_AC, d_BC)-biregular 2-CSP, Ψ', with alphabet size at most R, (d_A, d_B)-bounded degree 2-CSP, Ψ'', such that, with probability at least $2/3$,*

- *Completeness:* $\mathsf{val}(\Psi'') \geqslant \mathsf{val}(\Psi') - \nu_1$,
- *Soundness: If* $\mathsf{val}(\Psi'') \leqslant \frac{1}{R^{\nu_2}}$, *then* $\mathsf{val}(\Psi'') \leqslant \frac{1}{\nu_2 - \nu_1}\left(\frac{1}{d_A} + \frac{1}{d_B}\right)$

Putting everything together, we can prove Theorem 1.5.

Proof of Theorem 1.5. Recall the values η and d from Theorem 1.5. Start with an instance Ψ of 2-CSP from Theorem 7.1 with $\varphi = \varepsilon = 0.01\eta$ and sufficiently large alphabet size R. Then Ψ is (d_1, d_2)-biregular, with sufficiently large alphabet size R relative to $\varphi^{-1}, \varepsilon^{-1}$, and d. For such a Ψ, it is NP-hard to distinguish whether $\mathsf{val}(\Psi) = 1 - 0.01\eta$, or $\mathsf{val}(\Psi) \leqslant \frac{1}{R^{1-0.01\eta}}$.

Applying Lemma 7.6 with $c_1 = c_2 = d$ yields, in polynomial time, a (dd_1d_2, dd_1d_2)-biregular 2-CSP, Ψ', with alphabet size R and satisfying $\mathsf{val}(\Psi') = \mathsf{val}(\Psi)$. Next, applying Theorem 7.7, with $d_A = d, d_B = d, C = d_1d_2, \nu_1 = 0.01\eta, \nu_2 = 1 - \varepsilon$ to obtain, in randomized polynomial time, a 2-CSP, Ψ'', with degree at most d such that:

- If $\mathsf{val}(\Psi) \geqslant 1 - \varphi$, then $\mathsf{val}(\Psi'') \geqslant 1 - \varphi - \nu_1 = 1 - 0.02\eta$.
- If $\mathsf{val}(\Psi) \leqslant \frac{1}{R^{1-\varepsilon}}$, then $\mathsf{val}(\Psi'') \leqslant \frac{1}{1-\varepsilon-\nu_1}\left(\frac{1}{d} + \frac{1}{d}\right) = \frac{1}{1-0.02\eta} \cdot \frac{2}{d}$.

Finally note that,

$$\frac{1 - 0.02\varepsilon}{\frac{1}{1-0.02\varepsilon} \cdot \frac{2}{d}} \geqslant d\left(\frac{1}{2} - \varepsilon\right).$$

Thus, by Theorem 7.1 and the randomized polynomial time reduction above, it follows that unless $\mathsf{NP} = \mathsf{BPP}$, there is no polynomial time $d\left(\frac{1}{2} - \eta\right)$ approximation algorithm for 2-CSP with degree at most d. $\square$

7.4 Proof of Theorem 1.6

Combining our 2-Prover-1-Round Game in Theorem 1.3 with [Lae14, Lemma 4] and [Man19, Theorem 1], we obtain improved hardness of approximation results for Rooted k-connectivity on undirected graphs, the vertex-connectivity survivable network design problem, and the vertex-connectivity k-route cut problem on undirected graphs. The reduction is exactly the same as the reduction therein and we therefore omit the details.

8 Bounding the Number of Successful Zoom-outs of a Fixed Codimension

The goal of this section is to prove Lemma 5.19. Let $\mathbb{F}_q^n = \mathbb{F}_q^V$ be the space that we are working in and suppose T is a table assigning linear functions to $\mathsf{Grass}_q(n, 2\ell)$. We assume $n \geqslant 2^{100}q^\ell$. Let us review the set up of Lemma 5.19. Recall that we set

$$\xi = \delta^5, \quad \delta_2 = \xi/100, \quad t = 2^{2+10/\delta_2}.$$

Let $\mathcal{S} = \{W_1, \dots, W_N\}$ be a set of codimension r-subspaces in $\mathbb{F}_q^n$ of size

$$N \geqslant q^{100(t-1)!r^2\ell\xi^{-1}}.$$

For each W_i, let $f_i : W_i \to \mathbb{F}_q$ be a linear function such that $f_i|_L = T[L]$ for at least C-fraction of the 2ℓ-subspaces $L \in W_i$, where $C \geqslant q^{-2(1-\xi)\ell}$, and $\xi > 0$.

8.1 Step 1: Reducing to a Generic Set of Subspaces

Applying Lemma 5.8, with parameter t as defined, we get that there exists a subspace $V' \subseteq \mathbb{F}_q^n$ and a set of

$$m_1 \geqslant \frac{N^{\frac{1}{(r+1)\cdot(t-1)!}}}{q^r} \geqslant q^{75r\ell\xi^{-1}}$$

subspaces $\mathcal{W} = \{W_1, \dots, W_{m_1}\} \subseteq \mathcal{S}$, such that

- Each $W_i \in \mathcal{W}$ is contained in V' and has co-dimension s with respect to V', where $s \leqslant r$.
- $\mathcal{W}$ is t-generic with respect to V'.

We remark that this subspace V' will ultimately be the one used for Lemma 5.19. The remainder of the proof is devoted to finding the linear function $h' : V' \to \mathbb{F}_q$, and the set $\mathcal{W}'$, which will be a subset of $\mathcal{W}$ above.

8.2 Step 2: Finding Local Agreement

For a subspace X and linear assignment to X, $\sigma \in \mathbb{F}_q^X$, let

$$\mathcal{L}_X = \{L \in \mathcal{L} \mid X \subseteq L\} \quad \text{and} \quad \mathcal{L}_{X,\sigma} = \{L \in \mathcal{L}_X \mid T[L]|_X = \sigma\}.$$

Likewise, define

$$\mathcal{W}_X = \{W_i \in \mathcal{W} \mid X \subseteq W_i\} \quad \text{and} \quad \mathcal{W}_{X,\sigma} = \{W_i \in \mathcal{W}_X \mid f_i|_X = \sigma\}.$$

The first step of our proof is to find sets $\mathcal{W}_{X,\sigma}$ and $\mathcal{L}_{X,\sigma}$ that have strong agreement between them, in the sense of the following lemma. The approach of this first step is similar to that of [IKW12, BDN17, MZ23]. Fix $\gamma > 0$ to be a small constant, say $\gamma = 10^{-6}$.

Lemma 8.1. *There exists a $2\left(1-\frac{\xi}{2}\right)\ell$-dimensional subspace X, a linear assignment, σ, to X, and sets $\mathcal{L}_{X,\sigma}$ and $\mathcal{W}_{X,\sigma}$ such that the following hold:*

- $\mu_X(\mathcal{L}_{X,\sigma}) \geqslant \frac{C}{6}$.
- $|\mathcal{W}_{X,\sigma}| \geqslant \frac{m_1}{q^{10r\ell}}$.
- *Choosing $L \in \mathcal{L}_{X,\sigma}$ uniformly, and $W_i \in \mathcal{W}_{X,\sigma}$ uniformly such that $W_i \supseteq L$, we have*
$$\Pr_{L \subseteq W_i}[f_i|_L \neq T[L]] \leqslant 5\gamma.$$

Proof. Deferred to Appendix E □

As an immediate corollary, we have the following statement. The difference between Corollary 8.2 and Lemma 8.1 is that the former we require the third condition to hold every every $L \in \mathcal{L}_{X,\sigma}$, instead of a random L, and we also require every $L \in \mathcal{L}_{X,\sigma}$ to be contained in roughly the same number of $W_i \in \mathcal{W}_{X,\sigma}$.

Corollary 8.2. *Taking $\mathcal{L}_{X,\sigma}$ and $\mathcal{W}_{X,\sigma}$ from Lemma 8.1, there is a subset $\mathcal{L}'_{X,\sigma} \subseteq \mathcal{L}_{X,\sigma}$ such that the following hold.*

- $\mu_X(\mathcal{L}'_{X,\sigma}) \geqslant \frac{C}{12}$
- $m_2 = |\mathcal{W}_{X,\sigma}| \geqslant \frac{m_1}{q^{10r\ell}}$.
- *For every $L \in \mathcal{L}'_{X,\sigma}$, choosing $W_i \in \mathcal{W}_{X,\sigma}$ uniformly such that $W_i \supseteq L$, we have*
$$\Pr_{W_i \supseteq L, W_i \in \mathcal{W}_{X,\sigma}}[f_i|_L \neq T[L]] \leqslant 12\gamma.$$
- *For every $L \in \mathcal{L}'_{X,\sigma}$,*
$$0.95 \cdot |\mathcal{W}_{X,\sigma}| \cdot q^{-\xi\ell\cdot s} \leqslant N_{\mathcal{W}_{X,\sigma}}(L) \leqslant 1.05 \cdot |\mathcal{W}_{X,\sigma}| \cdot q^{-\xi\ell\cdot s},$$

where $N_{\mathcal{W}_{X,\sigma}}(L) = |\{W_i \supseteq L \mid W_i \in \mathcal{W}_{X,\sigma}\}|$.

Proof. Take $X, \sigma, \mathcal{L}_{X,\sigma}$, and $\mathcal{W}_{X,\sigma}$ as guaranteed by Lemma 8.1, so that $\mu_X(\mathcal{L}_{X,\sigma}) \geqslant \frac{C}{6}$. We will keep the same $X, \sigma,$, but we remove some L's from $\mathcal{L}_{X,\sigma}$ to make the third and fourth items hold.

By Markov's inequality, at most $\frac{5}{12}$-fraction of $L \in \mathcal{L}_{X,\sigma}$ violate the third item. By Lemma 5.16 applied to $\mathcal{W}_{X,\sigma}$ with parameters $Q = X$, $j = 2\ell$, $a = 2\left(1-\frac{\xi}{2}\right)\ell$, $r = s$, and $c = 0.05$, we have that

$$\Pr_{L \in \mathsf{Zoom}[Q,V]}\left[\left|N_{X,\sigma}(L) - q^{-\xi\ell\cdot s}m_2\right| \geqslant 0.1q^{-\xi\ell}m_2\right] \leqslant \frac{400q^{-\xi\ell\cdot s}}{m_2} \leqslant \frac{C}{100},$$

where $N_{X,\sigma}(L) = |\{W_i \supseteq L \mid W_i \in \mathcal{W}_{X,\sigma}\}|$. It follows that after removing the $L \in \mathcal{L}_{X,\sigma}$ that do not satisfy the third or fourth condition, we arrive at the desired $\mathcal{L}'_{X,\sigma}$, which still has measure at least

$$\frac{7}{12} \cdot \frac{C}{6} - \frac{C}{100} \geqslant \frac{C}{12}.$$ □

We fix X, σ as well as $W_{X,\sigma}$ and $\mathcal{L}'_{X,\sigma}$ as in Corollary 8.2 throughout the rest of the argument.

8.3 Step 3: A Global Set with Local Agreement

The next step is to further refine the set $\mathcal{L}'_{X,\sigma}$ so that the remaining subspaces "evenly cover" a subspace of $V^\star \subseteq V'$ with codimension $\dim(X) + O_{\delta_2}(1)$. To do this, we will reduce to the case where $\mathcal{L}'_{X,\sigma}$ is global within some zoom-in A and zoom-out B such that $X \subseteq A \subseteq B$. This is done via the following argument, which we outline below:

1. While $\mathcal{L}'_{X,\sigma}$ is not global, there must be some zoom-in or zoom-out on which it is dense, so consider the restriction to this zoom-in or zoom-out.
2. This increases the measure of $\mathcal{L}'_{X,\sigma}$, and we may repeat until we have a global set (within some zoom-in, zoom-out combination).
3. By choosing the globalness parameter suitably, we are able to perform the above process in relatively few times until the restriction of $\mathcal{L}'_{X,\sigma}$ that we arrive at is global.
4. As a result, when restricting to the zoom-in, zoom-out combination, the resulting set of subspaces evenly covers a space that is still relatively large in the sense that it contains $q^{-O(\xi^{-1})}$-fraction of V_0.

For a zoom-in A and zoom-out B such that $X \subseteq A \subseteq B \subseteq V'$, write $V' = A \oplus V_0$ and $B = A \oplus V^\star$, where $V^\star \subseteq V_0$. Also define

$$\mathcal{W}^\star_{[A,B]} = \{W_i^\star \mid \exists W_i \in \mathcal{W}_{X,\sigma} \text{ s.t } A \oplus W_i^\star = W_i \cap B\}.$$

It is clear that each $W_i^\star \in \mathcal{W}^\star_{[A,B]}$ is contained inside of some $W_i \in \mathcal{W}_{X,\sigma}$, so for each $W_i^\star$ we may define the restriction of f_i to $W_i^\star$ as $f_i^\star = f_i|_{W_i^\star}$.

Lemma 8.3. *There is a zoom-in A and a zoom-out B such that the following holds with the notation above. There exists a collection of subspaces $\mathcal{W}^\star = \{W_1^\star, \dots, W_{m_3}^\star\} \subseteq \mathcal{W}^\star_{[A,B]}$ of codimension s with respect to $V^\star$ such that:*

1. *For some $\ell' \geqslant \frac{\xi}{3}\ell$ there exists $\mathcal{L}^\star \subseteq \mathsf{Grass}_q(V^\star, \ell')$ such that $\mu(\mathcal{L}^\star) = \eta \geqslant \frac{C}{12}$.*
2. *The set $\mathcal{L}^\star$ is $(1, q^{\delta_2 \ell}\eta)$-pseudo-random.*
3. *Each $W_i^\star$ has codimension $s \leqslant r$ inside of $V^\star$ and $\mathcal{W}^\star$ is 4-generic, with respect to $V^\star$.*
4. *$m_3 \geqslant \frac{q^{-10s/\delta_2}}{2} \cdot m_2$.*
5. *For every $L \in \mathcal{L}^\star$, choosing $W_i^\star \in \mathcal{W}^\star$ uniformly such that $W_i^\star \supseteq L$, we have*
$$\Pr_{W_i^\star \supseteq L, W_i^\star \in \mathcal{W}^\star}[f_i|_L \neq T[L]] \leqslant 14\gamma.$$
6. *For every $L \in \mathcal{L}^\star$,*
$$0.8 \cdot m_3 \cdot q^{-s\cdot\ell'} \leqslant N_{\mathcal{W}^\star}(L) \leqslant 1.2 \cdot m_3 \cdot q^{-s\cdot\ell'},$$
where $N_{\mathcal{W}^\star}(L) = |\{W_i^\star \in \mathcal{W}^\star \mid W_i^\star \supseteq L\}|$

Here the table T is assigns linear functions to $L \in \mathsf{Grass}_q(V^\star, \ell)$, and is essentially the original table, i.e

$$T[L] \equiv T[A \oplus L]|_L.$$

Proof. Deferred to Appendix E.2. $\square$

Finally, as a consequence of pseudo-randomness, we may apply Lemma 2.4, to get that $\mathcal{L}^\star$ evenly covers $V^\star$.

Lemma 8.4. *Setting $Z = \{z \in V^\star \mid |\mu_z(\mathcal{L}^\star) - \eta| \leqslant \frac{\eta}{10}\}$, we have that,*

$$\frac{|Z|}{|V^\star|} \geqslant 1 - q^{\frac{\ell'}{2}}.$$

Proof. This is immediate by the pseudo-randomness of $\mathcal{L}^\star$ and Lemma 2.4. $\square$

8.4 Step 4: Local to Global Agreement

Lemma 8.5. *We have*

$$\Pr_{\substack{W_i^\star, W_j^\star \in \mathcal{W}^\star \\ z \in W_i^\star \cap W_j^\star \cap Z}} [f_i^\star(z) \neq f_j^\star(z)] \leqslant 500\gamma,$$

and for every $W_i^\star, W_j^\star \in \mathcal{W}^\star$,

$$|W_i^\star \cap W_j^\star \cap Z| \geqslant 0.81 \cdot |W_i^\star \cap W_j^\star|.$$

Proof. Deferred to Section E.3. $\square$

Using Lemma 8.5, we conclude the proof of Lemma 5.19 by using ideas from the Raz-Safra analysis of the Plane versus Plane test [RS97]. Define the following graph, G, with vertex set $\mathcal{W}^\star$ and an edge between $W_i^\star, W_j^\star$ if and only if $f_i^\star|_{W_i^\star \cap W_j^\star} = f_j^\star|_{W_i^\star \cap W_j^\star}$. We claim that this graph contains a large clique. To do so, we show that the graph is nearly transitive. For a graph $H = (V, E)$, define

$$\beta(H) = \max_{(u,w) \notin E} \Pr_v[(v, u), (v, w) \in E].$$

A graph H is transitive $\beta(H) = 0$. It is easy to see that transitive graphs are (edge) disjoint unions of cliques. The following lemma, proved in [RS97], asserts that if H is relatively dense and $\beta(H)$ is small, then one could remove only a small fraction of the edges and get a fully transitive graph.

Lemma 8.6. *[RS97, Lemma 2] Any graph $H = (V, E)$ can be made transitive by deleting at most $3\sqrt{\beta(H)}|V|^2$ edges.*

To use lemma 8.6 we first show that the graph G we defined is highly transitive.

Claim 8.7. *We have $\beta(G) \leqslant \frac{1}{m_3}$.*

Proof. Fix a $W_i^\star, W_j^\star$ that are not adjacent. We claim that they can have at most 1 common neighbor. Suppose for the sake of contradiction that $W_a^\star, W_b^\star$ are distinct common neighbors. Then,

$$f_i^\star|_{W_i^\star \cap W_j^\star \cap W_a^\star} = f_j^\star|_{W_i^\star \cap W_j^\star \cap W_a^\star},$$

and

$$f_i^\star|_{W_i^\star \cap W_j^\star \cap W_b^\star} = f_j^\star|_{W_i^\star \cap W_j^\star \cap W_b^\star}.$$

It follows that $f_i^\star$ and $f_j^\star$ agree on $W_i^\star \cap W_j^\star \cap W_a^\star \oplus W_i^\star \cap W_j^\star \cap W_b^\star$. However, since $\mathcal{W}^\star$ is 4-generic, we have

$$\mathrm{codim}(W_i^\star \cap W_j^\star \cap W_a^\star \oplus W_i^\star \cap W_j^\star \cap W_b^\star) \leqslant 3s + 3s - 4s = 2s$$

and

$$W_i^\star \cap W_j^\star \cap W_a^\star \oplus W_i^\star \cap W_j^\star \cap W_b^\star \subseteq W_i^\star \cap W_j^\star,$$

so it must be the case that $W_i^\star \cap W_j^\star \cap W_a^\star \oplus W_i^\star \cap W_j^\star \cap W_b^\star = W_i^\star \cap W_j^\star$. This contradicts the assumption that $W_i^\star$ and $W_j^\star$ are not adjacent. Thus, any two non-adjacent vertices can have at most 1 common neighbor, and the result follows. □

Claim 8.8. *The graph G contains a clique of size of size $\frac{m_3}{2}$*

Proof. Applying Markov's inequality and a union bound to Lemma 8.5, we have that with probability at least $9/10$ over $W_i^\star$ and $W_j^\star$, we have both $\Pr_{z \in W_i^\star \cap W_j^\star \cap Z}[f_i^\star(z) \neq f_j^\star(z)] \leqslant 10001\gamma$ and $|W_i^\star \cap W_j^\star \cap Z| \geqslant 0.81 \cdot |W_i^\star \cap W_j^\star|$. In this case, $f_i^\star$ and $f_j^\star$ agree on at least $(1 - 10001\gamma)$-fraction of the points in $W_i^\star \cap W_j^\star \cap Z$, which is in turn at least $(1 - 10001\gamma) \cdot 0.8 > 1/q$-fraction of the points in $W_i^\star \cap W_j^\star$. As $f_i^\star$ and $f_j^\star$ are linear functions, the Schwartz-Zippel lemma implies that such $W_i^\star, W_j^\star$ are adjacent in G and that G has at least $81m_3^2/100$ edges.

By Claim 8.7 and Lemma 8.6, we can delete $2m_3^{3/2}$ edges to make G a union of cliques. Doing so yields a graph on m_3 vertices with at least $m_3^2/2$-edges that is a union of cliques. Let $C_1, \ldots, C_N$ be the cliques, with C_1 being the largest one. We have,

$$|C_1| \cdot m \geqslant |C_1| \cdot \sum_{i=1}^{N} |C_i| \geqslant \sum_{i=1}^{N} |C_i|^2 \geqslant \frac{m_3^2}{2}.$$

It follows that $|C_1| \geqslant \frac{m_3}{2}$, and that G contains a clique of size at least $\frac{m_3}{2}$. □

Let $\mathcal{C}$ be the clique guaranteed by Claim 8.8 and write $\mathcal{C} = \{W_1^\star, \ldots, W_{\frac{m_3}{2}}^\star\}$. To complete the proof of Lemma 5.19, we will find a linear h, such that for all $1 \leqslant i \leqslant \frac{m_3}{2}$, $f_i^\star|_{V^\star} \equiv h|_{V^\star}$, and then show that this h can be extended to $X \oplus A \oplus V^\star$ in a manner that is consistent with many of the *original* f_i's for $1 \leqslant i \leqslant m_3/2$. To this end, first define $g : V^\star \to \mathbb{F}_q$ as follows:

$$g(x) = \begin{cases} f_i^\star(x), & \text{if } \exists W_i^\star \in \mathcal{C}, x \in W_i^\star \\ 0, & \text{otherwise.} \end{cases} \tag{15}$$

Since $f_i^\star(x) = f_j^\star(x)$ whenever $x \in W_i^\star \cap W_j^\star$, it does not matter which i is chosen if there are multiple $W_i^\star \in \mathcal{C}$ containing x. Thus, g is well defined and $g|_{W_i^\star} = f_i^\star$ for all $1 \leqslant i \leqslant \frac{m_3}{2}$.

We next show that g is close to a linear function and that this linear function agrees with most of the functions $f_i^\star|_{W_i^\star}$ for $W_i^\star \in \mathcal{C}$. To begin, we show that g passes the standard linearity test with high probability.

Lemma 8.9. *We have,*

$$\Pr_{z_1, z_2 \in \mathbb{F}_q^n}[g(z_1 + z_2) = g(z_1) + g(z_2)] \geqslant 1 - \frac{3q^{2s}}{m_3}.$$

Proof. Note that we have

$$\Pr_{z_1,z_2\in\mathbb{F}_q^n}[g(z_1+z_2)=g(z_1)+g(z_2)] \geqslant \Pr_{z_1,z_2\in\mathbb{F}_q^n}[\exists W_i^\star\in\mathcal{C},\text{ s.t. },z_1,z_2\in W_i^\star].$$

For every $z_1, z_2 \in \mathbb{F}_q^n$ linearly independent, we can let $N(z_1,z_2)$ denote the number of $W_i^\star \in \mathcal{C}$ containing $\mathrm{span}(z_1,z_2)$. The result them follows from Lemma 5.16 with $a=0$, $j=2$, $r=s$, and $c=0.99$. We have,

$$\begin{aligned}Pr_{z_1,z_2\in\mathbb{F}_q^n}[\exists W_i^\star\in\mathcal{C},\text{ s.t. },z_1,z_2\in W_i^\star] &\geqslant \Pr_{z_1,z_2\in\mathbb{F}_q^n}[N(z_1,z_2)>0 \mid \dim(\mathrm{span}(z_1,z_2))=2]-\frac{q}{q^n}\\ &\geqslant 1-\frac{3q^{3s}}{m_3},\end{aligned}$$

where $\frac{q}{q^n}$ is an upper bound on the probability that $\dim(\mathrm{span}(z_1,z_2))\neq 2$. □

Applying the linearity testing result of Blum, Luby, and Rubinfeld [BLR93, Theorem 4.1] we get that g is $\frac{12q^{2s}}{m_3}$-close to a linear function, say $h: V^\star\to\mathbb{F}_q$. We will conclude by showing that this h is the desired function which agrees with many of the original f_i's. To this end, we first show that agrees with many of the $f_i^\star$'s that we have (which are restrictions of the original f_i's), and then show that h can be extended to V' in a manner that retains agreement with many of the f_i's.

Towards the first step, set $S=\{x\in V^\star \mid g(x)\neq h(x)\}$. We show that choosing $W_i^\star\in\mathcal{C}$ randomly, and then a point $x\in W_i^\star$, it is unlikely that $x\in S$. Define the measure ν over nonzero points in $\mathbb{F}_q^n$ obtained by choosing $W_i\in\mathcal{C}$ uniformly at random and then $x\in W_i$ nonzero uniformly at random. Let μ be the uniform measure, so $\mu(S)\leqslant\frac{3q^{2s}}{m_3}$. Then $\nu(S)$ is precisely the probability of interest and can be upper bounded using Lemma 5.14, with parameters $a=0, j=1$, codimension s,

$$\nu(S)\leqslant\mu(S)+\frac{6q^{\frac{s}{2}}}{\sqrt{m_3}}\leqslant\frac{7q^{\frac{s}{2}}}{\sqrt{m_3}}. \tag{16}$$

Lemma 8.10. *We have $h|_{W_i^\star}\equiv f_i^\star$ for at least half of the $W_i^\star\in\mathcal{C}$.*

Proof. By Markov's inequality and Equation (16) with probability at least $1/2$, over $W_i^\star\in\mathcal{C}$, we have

$$\frac{|W_i^\star\cap S|}{|W_i^\star|}\leqslant\frac{14q^{\frac{s}{2}}}{\sqrt{m_3}}<1-\frac{1}{q},$$

and $f_i^\star$ and $h|_{W_i^\star}$ agree on more than $1/q$ of the points in $W_i^\star$. Since $f_i^\star$ and $h|_{W_i^\star}$ are both linear, by the Schwartz-Zippel Lemma that $h|_{W_i^\star}\equiv f_i^\star$, and the result follows. □

We are now ready to finish the proof of Lemma 5.19.

Proof of Lemma 5.19. Summarizing, we now have linear functions $f_i^\star: W_i^\star\to\mathbb{F}_q$ for $1\leqslant i\leqslant\frac{m_3}{4}$ and a linear function $h: V^\star\to\mathbb{F}_q$ such that $h|_{W_i^\star}=f_i^\star$. Furthermore, for each $f_i^\star, W_i^\star$, there is a f_i, W_i from Lemma 5.19such that $W_i\cap V^\star=W_i^\star$, $W_i\subseteq V'$, $f_i|_{W_i^\star}=f_i^\star$, and $f_i|_X=\sigma$.

Finally, we will extend h in a manner so that it agrees with many of these original functions f_i. To this end, recall that $V^\star$ satisfies,

$$A\oplus V^\star=B\subseteq V'.$$

and $\dim(A)+\operatorname{codim}(B)-\dim(X) \leqslant \frac{10}{\delta_2}$. Therefore, we may choose a random linear function $h' : V' \to \mathbb{F}_q$ conditioned on $h|_{V^\star} \equiv h$ and $h'|_X \equiv \sigma$. For any f_i, we have that

$$\Pr_{h'}[h'|_{W_i} \equiv f_i] \leqslant q^{-(\dim(A)+\operatorname{codim}(B)-\dim(X))} \leqslant q^{-\frac{10}{\delta_2}}.$$

Indeed there is a $q^{-(\dim(A)-\dim(X))}$ probability that $h|_A \equiv f_i|_A$, as we condition on $h'|_X \equiv \sigma \equiv f_i|_X$. Then, extending h' from B to V', there is at least a $q^{-\operatorname{codim}(B)}$ probability that h' is equal to f_i on these extra dimensions. It follows that there is a linear $h' : V' \to \mathbb{F}_q$ such that $h'|_{W_i} \equiv f_i$ for at least

$$\frac{m_3 q^{-\frac{10}{\delta_2}}}{4} \geqslant q^{50r\ell\xi^{-1}}$$

of the pairs f_i, W_i from Lemma 5.19. Take these W_i to be the set $\mathcal{W}'$ for Lemma 5.19. As these $\mathcal{W}' \subseteq \mathcal{W}$, they are 2-generic with respect to V' and have codimension $s \leqslant r$ in V'. □

References

[ABH+05] Sanjeev Arora, Eli Berger, Elad Hazan, Guy Kindler, and Muli Safra. On non-approximability for quadratic programs. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 206–215. IEEE, 2005.

[ALM+92] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and hardness of approximation problems. In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*, pages 14–23, 1992.

[AS97] Sanjeev Arora and Madhu Sudan. Improved low-degree testing and its applications. In *Proceedings of the twenty-ninth annual ACM Symposium on Theory of Computing*, pages 485–495, 1997.

[BDN17] Amey Bhangale, Irit Dinur, and Inbal Livni Navon. Cube vs. cube low degree test. In Christos H. Papadimitriou, editor, *8th Innovations in Theoretical Computer Science Conference, ITCS 2017, January 9-11, 2017, Berkeley, CA, USA*, volume 67 of *LIPIcs*, pages 40:1–40:31. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2017.

[Ber00] Piotr Berman. A d/2 approximation for maximum weight independent set in d-claw free graphs. *Nordic Journal of Computing*, 7(3):178–184, 2000.

[BKS19] Boaz Barak, Pravesh K. Kothari, and David Steurer. Small-set expansion in shortcode graph and the 2-to-2 conjecture. In Avrim Blum, editor, *10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA*, volume 124 of *LIPIcs*, pages 9:1–9:12. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2019.

[BL15] Abhishek Bhowmick and Shachar Lovett. The list decoding radius of reed-muller codes over small fields. In *Proceedings of the forty-seventh annual ACM symposium on Theory of Computing*, pages 277–285, 2015.

[BLR93] Manuel Blum, Michael Luby, and Ronitt Rubinfeld. Self-testing/correcting with applications to numerical problems. *J. Comput. Syst. Sci.*, 47(3):549–595, 1993.

[CCK08] Tanmoy Chakraborty, Julia Chuzhoy, and Sanjeev Khanna. Network design for vertex connectivity. In *Proceedings of the fortieth annual ACM symposium on Theory of computing*, pages 167–176, 2008.

[CGM13] Marek Cygan, Fabrizio Grandoni, and Monaldo Mastrolilli. How to sell hyperedges: The hypermatching assignment problem. In *Proceedings of the twenty-fourth annual ACM-SIAM symposium on Discrete algorithms*, pages 342–351. SIAM, 2013.

[Cha16] Siu On Chan. Approximation resistance from pairwise-independent subgroups. *Journal of the ACM (JACM)*, 63(3):1–32, 2016.

[CW04] Moses Charikar and Anthony Wirth. Maximizing quadratic programs: Extending grothendieck's inequality. In *45th Annual IEEE Symposium on Foundations of Computer Science*, pages 54–60. IEEE, 2004.

[DFRR23] Pavel Dvořák, Andreas Emil Feldmann, Ashutosh Rai, and Paweł Rzążewski. Parameterized inapproximability of independent set in h-free graphs. *Algorithmica*, 85(4):902–928, 2023.

[DKK+18] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In *Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing*, pages 376–389, 2018.

[DKK+21] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in grassmann graphs. *Israel Journal of Mathematics*, 243(1):377–420, 2021.

[EKL23a] David Ellis, Guy Kindler, and Noam Lifshitz. An analogue of bonami's lemma for functions on spaces of linear maps, and 2-2 games. In Barna Saha and Rocco A. Servedio, editors, *Proceedings of the 55th Annual ACM Symposium on Theory of Computing, STOC 2023, Orlando, FL, USA, June 20-23, 2023*, pages 656–660. ACM, 2023.

[EKL23b] Shai Evra, Guy Kindler, and Noam Lifshitz. Hypercontractivity for global functions on the general linear group over a finite field. 2023+.

[Fei98] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *Journal of the ACM (JACM)*, 45(4):634–652, 1998.

[FGL+91] Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. 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*, pages 2–12, 1991.

[Gop10] Parikshit Gopalan. A fourier-analytic approach to reed-muller decoding. In *2010 IEEE 51st Annual Symposium on Foundations of Computer Science*, pages 685–694. IEEE, 2010.

[GRS00] Oded Goldreich, Ronitt Rubinfeld, and Madhu Sudan. Learning polynomials with queries: The highly noisy case. *SIAM J. Discret. Math.*, 13(4):535–570, 2000.

[Has99] Johan Hastad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Math*, 182:105–142, 1999.

[Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*, 48(4):798–859, 2001.

[IKW12] Russell Impagliazzo, Valentine Kabanets, and Avi Wigderson. New direct-product testers and 2-query pcps. *SIAM J. Comput.*, 41(6):1722–1768, 2012.

[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.

[KP06] Subhash Khot and Ashok Kumar Ponnuswami. Better inapproximability results for maxclique, chromatic number and min-3lin-deletion. In Michele Bugliesi, Bart Preneel, Vladimiro Sassone, and Ingo Wegener, editors, *Automata, Languages and Programming, 33rd International Colloquium, ICALP 2006, Venice, Italy, July 10-14, 2006, Proceedings, Part I*, volume 4051 of *Lecture Notes in Computer Science*, pages 226–237. Springer, 2006.

[KR03] Subhash Khot and Oded Regev. Vertex cover might be hard to approximate to within 2-\varepsilon. In *18th Annual IEEE Conference on Computational Complexity (Complexity 2003), 7-10 July 2003, Aarhus, Denmark*, page 379. IEEE Computer Society, 2003.

[KS13] Subhash Khot and Muli Safra. A two-prover one-round game with strong soundness. *Theory Comput.*, 9:863–887, 2013.

[Lae14] Bundit Laekhanukit. 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*, pages 1626–1643. SIAM, 2014.

[LM23] Euiwoong Lee and Pasin Manurangsi. Hardness of approximating bounded-degree max 2-csp and independent set on k-claw-free graphs. *CoRR*, abs/2309.04099, 2023.

[Man19] Pasin Manurangsi. A note on degree vs gap of min-rep label cover and improved inapproximability for connectivity problems. *Inf. Process. Lett.*, 145:24–29, 2019.

[Meg01] Alexandre Megretski. 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*, pages 365–392. Springer, 2001.

[MR10] Dana Moshkovitz and Ran Raz. Two-query PCP with subconstant error. *J. ACM*, 57(5):29:1–29:29, 2010.

[MZ23] Dor Minzer and Kai Zheng. Approaching the soundness barrier: A near optimal analysis of the cube versus cube test. In Nikhil Bansal and Viswanath Nagarajan, editors, *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 2761–2776. SIAM, 2023.

[NRT99] Arkadi Nemirovski, Cornelis Roos, and Tamás Terlaky. On maximization of quadratic form over intersection of ellipsoids with common center. *Mathematical programming*, 86(3):463–473, 1999.

[Nut12] Zeev Nutov. Approximating minimum-cost connectivity problems via uncrossable bifamilies. *ACM Transactions on Algorithms (TALG)*, 9(1):1–16, 2012.

[Rao08] Anup Rao. Parallel repetition in projection games and a concentration bound. In *Proceedings of the Fortieth Annual ACM Symposium on Theory of Computing*, pages 1–10, 2008.

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.

[RS97] Ran Raz and Shmuel Safra. 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*, pages 475–484, 1997.

[Tre14] Luca Trevisan. Inapproximability of combinatorial optimization problems. *Paradigms of Combinatorial Optimization: Problems and New Approaches*, pages 381–434, 2014.

[TW23] Theophile Thiery and Justin Ward. An improved approximation for maximum weighted k-set packing. In *Proceedings of the 2023 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1138–1162. SIAM, 2023.

A Proofs of Lemmas 2.3 and 2.4

In this section we prove Lemmas 2.3 and 8.4. The proofs of these lemmas requires tools from [EKL23a, EKL23b] regarding Fourier analysis over the Bilinear Scheme.

A.1 Fourier Analysis over the Bilinear Scheme

The key to proving Lemma 2.3 is a level-d inequality for indicator functions on the Bilinear Scheme due to Evra, Kindler, and Lifshitz [EKL23b]. In order to use the result of [EKL23b], however, we first give some necessary background for Fourier analysis over the Bilinear Scheme, and describe the analogues of zoom-ins, zoom-outs, and pseudorandomness. The latter is done in [EKL23a, EKL23b]. After doing so, we must then find a suitable map from the Grassmann graph to the Bilinear Scheme that (1) preserves the edges of our original bipartite inclusion graph between 2ℓ-dimensional and $2(1-\delta)\ell$ subspaces, and (2) maps zoom-ins and zoom-outs in the Grassmann graph to their analogues over the Bilinear Scheme.

The Bilinear Scheme: Let $\mathbb{F}_q^{n\times 2\ell}$ be the set of $n\times 2\ell$ matrices over $\mathbb{F}_q$. One can define a graph over $\mathbb{F}_q^{n\times 2\ell}$ that is similar to the Grassmann graphs by calling $M_1, M_2 \in \mathbb{F}_q^{n\times 2\ell}$ adjacent if $\dim(\ker(M_1 - M_2)) \leqslant s$ for some $s \leqslant 2\ell$. A graph of this form are often referred to as the *Bilinear Scheme*. For our purposes, we will need to work with a bipartite version of this graph between $\mathbb{F}_q^{(n-2\ell)\times 2\ell}$ and $\mathbb{F}_q^{(n-2(1-\delta)\ell)\times 2(1-\delta)\ell}$.

We equip the space $L_2(\mathbb{F}_q^{n\times 2\ell})$ with the following inner product:

$$\langle F, G\rangle = \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}}[F(M)\overline{G(M)}],$$

where the distribution taken over M is uniform. Let ω be a primitive pth root of unity, where recall p is the characteristic of $\mathbb{F}_q$. For $s \in \mathbb{F}_q^n$ and $x \in \mathbb{F}_q^n$, let $\chi_s(x) = \omega^{\mathrm{Tr}(s\cdot x)}$ where $\mathrm{Tr}\colon \mathbb{F}_q \to \mathbb{F}_p$ is the trace map.

Then, the characters, $\chi_S : \mathbb{F}_q^{n\times 2\ell} \to \mathbb{C}$ over all $S = (s_1, \ldots, s_{2\ell}) \in \mathbb{F}_q^{n\times 2\ell}$, given by

$$\chi_S(x_1, \ldots, x_{2\ell}) = \prod_{i=1}^{2\ell} \chi_{s_i}(x_i) = \omega^{\sum_{i=1}^{2\ell} \mathrm{Tr}(s_i \cdot x_i)}$$

form an orthonormal basis of $L_2(\mathbb{F}_q^{n\times 2\ell})$. As a result, any $F \in L_2(\mathbb{F}_q^{n\times 2\ell})$ can be expressed as,

$$F = \sum_{S \in \mathbb{F}_q^{n\times 2\ell}} \widehat{F}(S)\chi_S,$$

where $\widehat{F}(S) = \langle F, \chi_S \rangle$. The level d component of F is given by

$$F^{=d} = \sum_{S\colon \mathrm{rank}(S)=d} \widehat{F}(S)\chi_S.$$

If a function F only consists of components up to level d, i.e. $\widehat{F}(S) = 0$ for all $\mathrm{rank}(S) > d$, then we say F is of degree d.

We now describe the analogues of zoom-ins and zoom-outs on $\mathbb{F}_q^{n\times 2\ell}$. We also define the analogous notion of (r, ε)-pseudo-randomness for Boolean functions over $\mathbb{F}_q^{n\times 2\ell}$, and we begin by defining the analog of zoom-ins.

Definition A.1. *A zoom-in of dimension d over $\mathbb{F}_q^{n\times 2\ell}$ is given by d-pairs of vectors $(u_1, v_1), \ldots, (u_r, v_r)$ where each $u_i \in \mathbb{F}_q^{2\ell}$ and each $v_i \in \mathbb{F}_q^{n}$. Let $U \in \mathbb{F}_q^{2\ell\times d}$ and $V \in \mathbb{F}_q^{n\times d}$ denote the matrices whose ith columns are u_i and v_i respectively. Then the zoom-in on (U, V) is the set of $M \in \mathbb{F}_q^{n\times 2\ell}$ such that $MU = V$, or equivalently, $Mu_i = v_i$ for $1 \leqslant i \leqslant d$.*

Next, we define the analog of zoom-outs.

Definition A.2. *A zoom-out of dimension d is defined similarly, except by multiplication on the left. Given $X \in \mathbb{F}_q^{d\times n}$ and $Y \in \mathbb{F}_q^{d\times 2\ell}$, whose rows are given by x_i and y_i respectively, the zoom-out (X, Y) is the $M \in \mathbb{F}_q^{n\times 2\ell}$ such that $XM = Y$, or equivalently, $x_iM = y_i$ for $1 \leqslant i \leqslant d$.*

Let $\mathsf{Zoom}[(U, V), (X, Y)]$ denote the intersections of the zoom-in on (U, V) and the zoom-out on (X, Y). The codimension of $\mathsf{Zoom}[(U, V), (X, Y)]$ is the sum of the number of columns of U and the number of rows of X, which we will denote by $\dim(U)$ and $\mathrm{codim}(X)$. For a zoom-in and zoom-out pair and a Boolean function F, we define $F_{(U,V),(X,Y)} : \mathsf{Zoom}[(U, V), (X, Y)] \to \{0, 1\}$ to be the restriction of F which is given as

$$F_{(U,V),(X,Y)}(M) = F(M) \quad \text{for} \quad M \in \mathsf{Zoom}[(U, V), (X, Y)].$$

When $\dim(U) + \mathrm{codim}(X) = d$, we say that the restriction is of size d. We define (d, ε)-pseudo-randomness in terms of the L_2-norms of restrictions of F of size d. Here and throughout, when we consider restricted functions, the underlying measure is the uniform measure over the corresponding zoom-in and zoom-out set $\mathsf{Zoom}[(U, V), (X, Y)]$.

Definition A.3. *We say that an indicator function $F \in L_2(\mathbb{F}_q^{n\times 2\ell})$ is (d, ε)*-pseudorandom *if for all zoom-in zoom-out combinations $\mathsf{Zoom}[(U, V), (X, Y)]$ such that $\dim(U) + \mathrm{codim}(X) = d$, we have*

$$\left\| F_{(U,V),(X,Y)} \right\|_2^2 \leqslant \varepsilon.$$

We note that for Boolean functions F, $\left\|F_{(U,V),(X,Y)}\right\|_2^2 = \mathbb{E}_{M\in \mathsf{Zoom}[(U,V),(X,Y)]}[F(M)]$, and hence the definition above generalizes the definition we have for Boolean functions.

Definition A.4. *We say that an indicator function $F \in L_2(\mathbb{F}_q^{n\times 2\ell})$ is (d,ε,t)-*pseudo-random *if for all* $\mathsf{Zoom}[(U,V),(X,Y)]$ *such that* $\dim(U)+\operatorname{codim}(X) = d$*, we have,*

$$\left\|F_{(U,V),(X,Y)}\right\|_{t/(t-1)} = \left(\mathop{\mathbb{E}}_{M\in \mathsf{Zoom}[(U,V),(X,Y)]}\left[|F(M)|^{\frac{t}{t-1}}\right]\right)^{\frac{t-1}{t}} \leqslant \varepsilon.$$

The following result is a combination of two results form [EKL23b]. Roughly speaking, it states that if a Boolean function F is (r,ε) pseudo-random, then its degree d parts are $(r, C_{q,d}\varepsilon^2)$ pseudo-random for $d\leqslant r$.

Lemma A.5. *[EKL23b, Theorem 5.5 + Lemma 3.6] Let $t\geqslant 4$ be a power of 2 and let $F : \mathbb{F}_q^{n\times 2\ell} \to \{0,1\}$ be a function that is (d,ε,t)-pseudo-random. Then $F^{=d}$ is $(r, q^{10dr+500d^2t}\varepsilon^2)$-pseudo-random for all $r\geqslant d$.*

Proof. This lemma does not actually appear in [EKL23b], but it is easy to derive by combining Theorem 5.5 with Lemma 3.6 therein. In [EKL23a, EKL23b], the authors introduce an additional notion of generalized influences and having small generalized influences. We refrain from defining these notions explicitly as it is slightly cumbersome, but roughly speaking, one defines a Laplacian for each zoom-in zoom-out combination, and having (d,ε) small generalized influences means that applying these Laplacians on F, the 2-norm squared of the resulting function never exceeds ε.

With this notion in hand, if a function F is (d,ε,t)-pseudo-random, then by [EKL23b, Theorem 5.5] we get that $F^{=d}$ has $(d, q^{500d^2t}\varepsilon^2)$-small generalized influences. Applying [EKL23b, Lemma 3.6] then implies that $F^{=d}$ is $(r, q^{10dr}\cdot q^{500d^2t}\varepsilon^2)$-pseudo-random for any $r\geqslant d$, which is the desired result □

Lastly, we need the following global hypercontrativity result also due to [EKL23b].[9]

Theorem A.6. *[EKL23b, Theorem 1.11] Let $t\geqslant 4$ be a power of 2 and let $F : \mathbb{F}_q^{n\times 2\ell} \to$ be a function of degree d that is (d,ε)-pseudo-random. Then,*

$$\|F\|_t^t \leqslant q^{200d^2t^2}\varepsilon^{t/2}.$$

Combining Lemma A.5 and Theorem A.6, we arrive at the following result which bounds the t-norm of the level d component of pseudo-random indicator functions. This result will be the key to showing an analogue of Lemma 2.3 over the Bilinear Scheme.

Theorem A.7. *Let $t\geqslant 4$ be a power of 2. Then if $F : \mathbb{F}_q^{n\times 2\ell} \to \{0,1\}$ is (r,ε)-pseudo-random, we have*

$$\left\|F^{=d}\right\|_t \leqslant q^{500d^2t}\varepsilon^{\frac{t-1}{t}}$$

for all $d\leqslant r$.

[9] We remark that earlier results [EKL23a] showed similar statement for 4-norms, i.e. the case that $t=4$, and the result below follows by a form of induction on t. That is, one starts with F and concludes via applying the case $t=4$ that the function and it F^2 is $(d, C_{q,d}\varepsilon)$ pseudo-random. Then one apply the case $t=4$ on F^2 to conclude that F^4 is $(d, C'_{q,d}\varepsilon)$ pseudo-random and so on.

Proof. Suppose F is (r,ε)-pseudo-random, let $t \geqslant 4$ be a power of 2, and fix a $d \leqslant r$. Since $d \leqslant r$, we also have that F is (d,ε)-peudorandom. Therefore for any size d restriction of F, $F_{(U,V),(X,Y)}$,we have,

$$\left\|F_{(U,V),(X,Y)}\right\|_{t/(t-1)} = \left(\left\|F_{(U,V),(X,Y)}\right\|_2^2\right)^{\frac{t-1}{t}} \leqslant \varepsilon^{\frac{t-1}{t}}.$$

Thus, F is $(d,\varepsilon^{\frac{t-1}{t}},t)$-pseudo-random, and by Lemma A.5 it follows that $F^{=d}$ is $(d, q^{10d^2+500d^2t}\varepsilon^{\frac{2t-2}{t}})$-pseudo-random. Clearly, $F^{=d}$ is degree d, so applying Theorem A.6 we get that,

$$\left\|F^{=d}\right\|_t^t \leqslant q^{200d^2t^2}\left(q^{10d^2+500d^2t}\varepsilon^{\frac{2t-2}{t}}\right)^{t/2} \leqslant q^{500d^2t^2}\varepsilon^{t-1},$$

and taking t-th root finished the proof. □

A.2 An Analog of Lemmas 2.3 for the Bilinear Scheme

With Theorem A.7 in hand we can show an analogue of Lemma 2.3 for basis invariant functions over the Bilinear Scheme. To do so, we first define what we mean by basis invariant functions, then present an analogue of the adjacency operator $\mathcal{T}$ (which is originally defined for functions over subspaces) over the Bilinear Scheme, which we denote by $\mathcal{T}'$, and finally show that the previously described characters are eigenoperators of $\mathcal{T}'^* \circ \mathcal{T}'$, where $\mathcal{T}'^*$ is the adjoint of $\mathcal{T}$.

For a function $F \in L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$, say that F is basis invariant if $F(M) = F(MA)$ for any full rank $A \in \mathbb{F}_q^{2\ell\times 2\ell}$. We first show that the level d component of a basis invariant function is also basis invariant. The following identity regarding the characters will be useful.

Lemma A.8. *For any $S = (s_1,\dots,s_{\ell'}) \in \mathbb{F}_q^{n\times\ell'}$, any $M \in \mathbb{F}_q^{n\times 2\ell}$, and any matrix $A \in \mathbb{F}_q^{2\ell\times\ell'}$ we have,*

$$\chi_S(MA) = \chi_{SA^T}(M).$$

Proof. Letting $v_1,\dots,v_{2\ell}$ denote the columns of M and $a_{i,j}$ denote the entries of A, we have,

$$\chi_S(MA) = \omega^{\mathrm{Tr}\sum_{i=1}^{\ell'} s_i\cdot\left(\sum_{j=1}^{2\ell} v_j a_{j,i}\right)} = \omega^{\sum_{i=1}^{\ell'}\mathrm{Tr}\left(v_i\cdot\left(\sum_{j=1}^{2\ell} s_j a_{i,j}\right)\right)} = \chi_{SA^T}(M).$$ □

Lemma A.9. *Let $S = [s_1,\dots,s_{2\ell}] \in \mathbb{F}_q^{n\times 2\ell}$ and let $F \in L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$ be basis invariant. Then for any $A \in \mathbb{F}_q^{2\ell\times 2\ell}$ that is full rank, we have $\widehat{F}(SA) = \widehat{F}(S)$.*

Proof. For any matrix full rank $B \in \mathbb{F}_q^{2\ell\times 2\ell}$ we have

$$\begin{aligned}
\widehat{F}(S) &= \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}}\left[\chi_S(M)F(M)\right] \\
&= \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}}\left[\chi_S(M)F(MB^{-1})\right] \\
&= \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}}\left[\chi_S(MB)F(M)\right] \\
&= E_{M\in\mathbb{F}_q^{n\times 2\ell}}[\chi_{SB^T}(M)F(M)] \\
&= \widehat{F}(SB^T),
\end{aligned}$$

where we use that F is basis invariant in the third transition and Lemma A.8 in the fourth transition. Setting $B = A^T$ gives the result. □

Using Lemma A.9, we can show that the level d component of a basis invariant function is also basis invariant.

Lemma A.10. *If $F \in L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$ is basis invariant, then $F^{=d}$ is basis invariant as well for any d.*

Proof. Fix any $M \in \mathbb{F}_q^{n\times 2\ell}$ and $A \in \mathbb{F}_q^{2\ell\times 2\ell}$ full rank. We have

$$\begin{aligned}
F^{=d}(MA) &= \sum_{S\in\mathbb{F}_q^{n\times 2\ell},\mathrm{rank}(S)=d} \widehat{F}(S)\chi_S(MA) \\
&= \sum_{S\in\mathbb{F}_q^{n\times 2\ell},\mathrm{rank}(S)=d} \widehat{F}(S)\chi_{SA^T}(M) \\
&= \sum_{S\in\mathbb{F}_q^{n\times 2\ell},\mathrm{rank}(S)=d} \widehat{F}(S(A^T)^{-1})\chi_S(M) \\
&= \sum_{S\in\mathbb{F}_q^{n\times 2\ell},\mathrm{rank}(S)=d} \widehat{F}(S)\chi_S(M) \\
&= F^{=d}(M),
\end{aligned}$$

where we use Lemma A.8 in the second transition, and Lemma A.9 in the fourth transition. $\square$

We now define the following two operators which will be the analogues of $\mathcal{T}$ and $\mathcal{T}^*$ over the bilinear scheme. The first is $\mathcal{T}' : L_2\left(\mathbb{F}_q^{n\times 2\ell}\right) \to L_2\left(\mathbb{F}_q^{n\times 2(1-\delta)\ell}\right)$, given by:

$$\mathcal{T}'F(M') = \mathop{\mathbb{E}}_{v_1,\dots,v_{\delta\ell}}[F\left([M', v_1, \dots, v_{2\delta\ell}\right)].$$

In words, the operator $\mathcal{T}'$ averages over extensions of the matrix M to an $n \times 2\ell$ matrix by adding to it $2\delta\ell$ random columns. The next is $\mathcal{T}'^* : L_2\left(\mathbb{F}_q^{n\times 2(1-\delta)\ell}\right) \to L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$ given by:

$$\mathcal{T}'^*G(M) = \mathop{\mathbb{E}}_{A\in\mathbb{F}_q^{2\ell\times 2(1-\delta)\ell}}[G(MA) \mid \mathrm{rank}(A) = 2(1-\delta)\ell].$$

Strictly speaking, $\mathcal{T}'^*$ is not the adjoint of $\mathcal{T}'$; however, for the case where F is basis invariant, $\mathcal{T}'^*$ acts as the adjoint of $\mathcal{T}'$ in the following sense.

Lemma A.11. *For $F \in L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$ that is basis invariant and $G \in L_2\left(\mathbb{F}_q^{n\times 2(1-\delta)\ell}\right)$, we have*

$$\langle \mathcal{T}'F, G\rangle = \langle F, \mathcal{T}'^*G\rangle.$$

Proof. Let $J \in \mathbb{F}_q^{2\ell\times 2(1-\delta)\ell}$ be the matrix whose restriction to the first $2(1-\delta)\ell$ rows is the identity matrix $I_{2(1-\delta)\ell\times 2(1-\delta)\ell}$ and whose remaining rows are all 0. We have

$$\begin{aligned}
\langle \mathcal{T}'F, G\rangle &= \mathop{\mathbb{E}}_{M'\in\mathbb{F}_q^{n\times 2(1-\delta)\ell},v_1,\dots,v_{2\delta\ell}\in\mathbb{F}_q^n}\left[F'\left([M', v_1, \dots, v_{2\delta\ell}]\right)\cdot \overline{G(M')}\right] \\
&= \mathop{\mathbb{E}}_{M'\in\mathbb{F}_q^{n\times 2(1-\delta)\ell},v_i\in\mathbb{F}_q^n,A\in\mathbb{F}_q^{2\ell\times 2\ell}}\left[F'\left([M', v_1, \dots, v_{2\delta\ell}]A\right)\cdot \overline{G(M')} \mid \mathrm{rank}(A) = 2\ell\right] \\
&= \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell},A\in\mathbb{F}_q^{2\ell\times 2\ell}}[F(M)\cdot \overline{G(MA^{-1}J)} \mid \mathrm{rank}(A) = 2\ell],
\end{aligned}$$

where in the second transition we used the fact that F' is basis invariant, and in the third one we made a change of variables $M = [M', v_1, \ldots, v_{2\delta\ell}]A$. Now note that $A^{-1}J$ is the matrix A^{-1} restricted to its first $2(1-\delta)\ell$ columns and hence in the final distribution, $A^{-1}J$ is a uniformly random matrix in $\mathbb{F}_q^{n\times 2(1-\delta)\ell}$ with rank $2(1-\delta)\ell$. It follows that,

$$\langle \mathcal{T}'F, G\rangle = \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}, A\in\mathbb{F}_q^{2\ell\times 2(1-\delta)\ell}}[F(M)\cdot\overline{G(MA)} \mid \operatorname{rank}(A) = 2(1-\delta)\ell] = \langle F, \mathcal{T}'^{*}G\rangle. \qquad \square$$

We will want to understand the operator $\mathcal{T}^{*}\mathcal{T}'$, and towards this end we define the operator

$$\Phi F(M) = \mathop{\mathbb{E}}_{\substack{B\in\mathbb{F}_q^{n\times 2\delta\ell}\\ C\in\mathbb{F}_q^{2\delta\ell\times 2\ell}\\ \operatorname{rank}(C)=2\delta\ell}}[F(M+BC)].$$

The reason for introducing Φ is that, as the following lemma shows, it acts the same as $\mathcal{T}'^{*}\mathcal{T}'$ on basis invariant functions, but is easier to work with. This is due to the reason it is an averaging operator with respect to some Cayley graph over $\mathbb{F}_q^{n\times 2\ell}$, and therefore each character χ_S is an eigenvector of Φ and the eigenvalues have an explicit formula. These facts are shown in the next two lemmas respectively.

Lemma A.12. *If $F \in L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$ is basis invariant, then $\mathcal{T}'^{*}\mathcal{T}'F = \Phi F$.*

Proof. By definitions

$$\mathcal{T}'^{*}\mathcal{T}'F(M) = \mathop{\mathbb{E}}_{\substack{R'\in\mathbb{F}_q^{2\ell\times 2(1-\delta)\ell},\\ v_1,\ldots,v_{2\delta\ell}\in\mathbb{F}_q^{n}}}[F'\left([MR', v_1, \ldots, v_{2\delta\ell}] \mid \operatorname{rank}(R') = 2(1-\delta)\ell\right)].$$

We can also view $M' = [MR', v_1, \ldots, v_{2\delta\ell}]$ as being sampled as follows. Choose $R' \in \mathbb{F}_q^{2\ell\times 2(1-\delta)\ell}$ with linearly independent columns, extend R' to a matrix $R \in \mathbb{F}_q^{2\ell\times 2\ell}$ with linearly independent columns randomly by adding $2\delta\ell$ columns on the right, sample a random matrix $[0, \ldots, 0, w_1, \ldots, w_{2\delta\ell}] \in \mathbb{F}_q^{n\times 2\ell}$, and output,

$$M' = MR + [0, \ldots, 0, w_1, \ldots, w_{2\delta\ell}].$$

Furthermore, under this distribution, it is clear that $R \in \mathbb{F}_q^{2\ell\times 2\ell}$ is a uniformly random matrix with linearly independent columns. Therefore,

$$\begin{aligned}
\mathcal{T}'^{*}\mathcal{T}'F(M) &= \mathop{\mathbb{E}}_{\substack{R\in\mathbb{F}_q^{2\ell\times 2\ell},\\ w_1,\ldots,w_{2\delta\ell}\in\mathbb{F}_q^{n}}}[F(MR + [0, \ldots, 0, w_1, \ldots, w_{2\delta\ell}]]\\
&= \mathop{\mathbb{E}}_{\substack{R\in\mathbb{F}_q^{2\ell\times 2\ell},\\ w_1,\ldots,w_{2\delta\ell}\in\mathbb{F}_q^{n}}}[F(M + [0, \ldots, 0, w_1, \ldots, w_{2\delta\ell}]R^{-1})],
\end{aligned}$$

where we are using the fact that F is basis invariant and R is invertible. In the last expectation, note that the distribution over $[0, \ldots, 0, w_1, \ldots, w_{2\delta\ell}]R^{-1}$ is the same as that over BC where $B \in \mathbb{F}_q^{n\times 2\delta\ell}$ is uniformly random, and $C \in \mathbb{F}_q^{2\delta\ell\times 2\ell}$ is uniformly random conditioned on having linearly independent rows. More precisely, it is equal to BC where $B = [w_1, \ldots, w_{2\delta\ell}]$, and C is the last $2\delta\ell$ rows of R^{-1}. It follows that

$$\mathcal{T}'^{*}\mathcal{T}'F(M) = \mathop{\mathbb{E}}_{B\in\mathbb{F}_q^{n\times 2\delta\ell}, C\in\mathbb{F}_q^{2\delta\ell\times 2\ell}}[F(M+BC) \mid \operatorname{rank}(C) = 2\delta\ell]. \qquad \square$$

The following lemma gives upper bound on the eigenvalues of Φ.

Lemma A.13. *Suppose that* $\text{rank}(S) = t$. *If* $t = 0$, *then* χ_S *is an eigenvector of* Φ *of eigenvalue* 1. *Otherwise, if* $t > 0$, χ_S *is an eigenvector of* Φ *of eigenvalue which is at most* $3q^{t-n} + q^{-t(2\delta\ell-1)}$ *in absolute value.*

Proof. Fix S. We argued earlier that χ_S is an eigenvector of Φ, and we denote the corresponding eigevalue by $\lambda = \Phi\chi_S(0)$. If $t = 0$ the statement is clear, so we assume that $t > 0$ henceforth.

Find $A \in \mathbb{F}_q^{2\ell\times 2\ell}$ of full rank so that $SA^T = (v_1, \ldots, v_t, 0, 0, \ldots, 0)$ where $v_1, \ldots, v_t$ are linearly independent. Thus, as the distribution of C is invariant under multiplying by A^T from the right, we get that

$$\lambda = \Phi\chi_S(0) = \mathop{\mathbb{E}}_{B,C}[\chi_S(BCA^t) \mid \text{rank}(C) = 2\delta\ell] = \mathop{\mathbb{E}}_{B,C}[\chi_{SA}(BC) \mid \text{rank}(C) = 2\delta\ell],$$

and we may assume that $S = (v_1, \ldots, v_t, 0, \ldots, 0)$ for linearly independent $v_1, \ldots, v_t$ to begin with. Applying symmetry again, we conclude that

$$\lambda = \mathop{\mathbb{E}}_{\substack{v_1,\ldots,v_t\\ \text{linearly independent}}}\left[\Phi\chi_{(v_1,\ldots,v_t,\vec{0})}(0)\right] = \mathop{\mathbb{E}}_{\substack{v_1,\ldots,v_t\\ \text{linearly independent}}}\left[\mathop{\mathbb{E}}_{B,C} \omega^{\sum_{i=1}^{t}\text{Tr}(v_i\cdot\mathsf{col}_i(BC))}\right],$$

and interchanging the order of expectations we get that

$$\lambda = \mathop{\mathbb{E}}_{B,C}\left[\mathop{\mathbb{E}}_{\substack{v_1,\ldots,v_t\\ \text{linearly independent}}} \omega^{\sum_{i=1}^{t}\text{Tr}(v_i\cdot\mathsf{col}_i(BC))}\right],$$

Denote $w_i = \mathsf{col}_i(BC)$, and inspect these vectors.

Claim A.14. *If* $w_i \neq 0$ *for some* i, *then*

$$\left|\mathop{\mathbb{E}}_{\substack{v_1,\ldots,v_t\\ \textit{linearly independent}}}\left[\omega^{\sum_{i=1}^{t}\text{Tr}(v_i\cdot\mathsf{col}_i(BC))}\right]\right| \leqslant 2q^{t-n}.$$

Proof. We first claim that if $v_1, \ldots, v_t$ are chosen uniformly, then the left hand side is 0, or equivalently

$$\mathop{\mathbb{E}}_{\substack{v_1,\ldots,v_t\\ \text{uniform}}}\left[\omega^{\text{Tr}\left(\sum_{i=1}^{t} v_i\cdot w_i\right)}\right] = 0.$$

To see this, it suffices to show that $\sum_{i=1}^{t} v_i \cdot w_i$ takes every value in $\mathbb{F}_q$ with equal probability, and we focus on showing this. Fix i such that $w_i \neq 0$ and suppose the jth entry, $w_{i,j}$ is nonzero. We can fix all entries of the $v_1, \ldots, v_t$ uniformly except for $v_{i,j}$, and then for each $\alpha \in \mathbb{F}_q$, there is exactly one choice of $v_{i,j}$ that will result in $\sum_{i=1}^{t} v_i \cdot w_i = \alpha$.

Thus, if we took the distribution over $v_1, \ldots, v_t$ to be uniformly and independently chosen, then the magnitude of the above expectation would be 0. Hence, we conclude that the above expectation is at most twice the probability randomly chosen $v_1, \ldots, v_t$ are not linearly independent, which is at most q^{t-n}. □

By Claim A.14 we conclude that $\lambda \leqslant 2q^{t-n} + \Pr_{B,C}[w_i = 0\ \forall i = 1, \ldots, t]$, and we next bound this probability. Recalling the definition of w_i, we have that

$$w_i = \sum_{j=1}^{2\delta\ell} C(j,i)\mathsf{col}_j(B).$$

Consider the $2\delta\ell \times t$ minor of C and call it C'. First we upper bound the probability that $\operatorname{rank}(C') = 0$. Note that the distribution of C is the same as of $A|_{2\delta\ell\times 2\ell}$ where $A \in \mathbb{F}_q^{2\ell\times 2\ell}$ is a random invertible matrix. Thus, C' has the same distribution as of $A|_{2\delta\ell\times t}$, and the probability that $C' = 0$ is at most

$$\frac{q^{2\ell-2\delta\ell}}{q^{2\ell}-1}\cdot\frac{q^{2\ell-2\delta\ell}}{q^{2\ell}-q}\cdots\frac{q^{2\ell-2\delta\ell}}{q^{2\ell}-q^{t-1}} \leqslant q^{-t(2\delta\ell-1)}.$$

It remains to bound the probability that w_i are all 0 in the case that $\operatorname{rank}(C') \geqslant 1$. In this case, assume without loss of generality that the first column of C' is non-zero. Thus, it follows that over the randomness of B, the vector w_1 is uniformly chosen from $\mathbb{F}_q^n$, and so the probability it is the all 0 vector is at most q^{-n}. Combining, we get that $\lambda \leqslant 3q^{t-n} + q^{-t(2\delta\ell-1)}$. □

Finally, using Lemma A.8 again, we can show that $\mathcal{T}'^*$ does not increase the level of a function,

$$\mathcal{T}'^*\chi_S(M) = \mathop{\mathbb{E}}_{A\in\mathbb{F}_q^{2\ell\times 2(1-\delta)\ell}}[\chi_S(MA)] = \mathop{\mathbb{E}}_{A}[\chi_{SA^T}(M)], \tag{17}$$

and obtain a useful identity for decomposing inner products.

Lemma A.15. *Let $F \in L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$ and $G \in L_2\left(\mathbb{F}_q^{n\times 2(1-\delta)\ell}\right)$. Then,*

$$\langle \mathcal{T}' F^{=d}, G\rangle = \langle \mathcal{T}' F^{=d}, G^{=d}\rangle.$$

As a consequence we also have

$$\langle F^{=d}, \mathcal{T}'^* G\rangle = \langle \mathcal{T}' F^{=d}, \mathcal{T}'^* G^{=d}\rangle.$$

Proof. Using Equation (17), we have,

$$\mathcal{T}'^* G^{=j}(M) = \sum_{S\in\mathbb{F}_q^{n\times 2(1-\delta)\ell},\operatorname{rank}(S)=j} \widehat{G}(S)\mathop{\mathbb{E}}_{A}\left[\chi_{SA^T}(M) \mid \operatorname{rank}(A) = 2(1-\delta)\ell\right].$$

Since $\operatorname{rank}(S) = j$, it follows that the $\operatorname{rank}(SA^T)$ is at most j, so it follows that for $j < d$, we have $\langle F^{=d}, \mathcal{T}^* G^{=j}\rangle = 0$. As a result,

$$\langle \mathcal{T}' F^{=d}, G\rangle = \langle F^{=d}, \mathcal{T}'^* G\rangle = \sum_{j=d}^{2(1-\delta)\ell} \langle F^{=d}, \mathcal{T}'^* G\rangle. \tag{18}$$

Next we have,

$$\mathcal{T}' F^{=d}(M) = \sum_{S\in\mathbb{F}_q^{n\times 2\ell},\operatorname{rank}(S)=d} \widehat{F}(S)\chi_S(M') = \sum_{S\in\mathbb{F}_q^{n\times 2\ell},\operatorname{rank}(S)=d} \widehat{F}(S)\chi_{S'}(M'),$$

where both M' and S' are obtained from M by removing the last $2\delta\ell$ columns. It follows that $\mathcal{T}'F^{=d}$ has level at most d, so using (18) we get

$$\langle \mathcal{T}'F^{=d}, G\rangle = \sum_{j=d}^{2(1-\delta)\ell} \langle F^{=d}, \mathcal{T}'^{*}G^{=j}\rangle = \sum_{j=d}^{2(1-\delta)\ell} \langle \mathcal{T}'F^{=d}, G^{=j}\rangle = \langle \mathcal{T}F^{=d}, G^{=d}\rangle. \qquad \square$$

We are now ready to state and prove an analog of Lemma 2.3 for basis invariant functions on the Bilinear scheme.

Lemma A.16. *Let $F \in L_2(\mathbb{F}_q^{n\times 2\ell})$ and $G \in L_2(\mathbb{F}_q^{n\times 2(1-\delta)\ell})$ be basis invariant indicator functions with $\mathbb{E}[F] = \alpha, \mathbb{E}[G] = \beta$. If F is (r,ε) pseudo-random and basis invariant, then for all $t \geqslant 4$ that are powers of 2, we have*

$$\langle \mathcal{T}'F, G\rangle \leqslant q^{O_{t,r}(1)}\beta^{(t-1)/t}\varepsilon^{2t/(2t-1)} + q^{-r\delta\ell}\sqrt{\alpha\beta}.$$

Proof. Using the degree decomposition of F and Lemma A.15, we can write

$$\langle \mathcal{T}'F, G\rangle = \sum_{d=0}^{2\ell}\langle \mathcal{T}'F^{=d}, G^{=d}\rangle.$$

We first bound the contribution from terms in the summation with $d > r$ using Cauchy-Schwarz. For $d > r$,

$$\begin{aligned}
|\langle \mathcal{T}'F^{=d}, G^{=d}\rangle|^2 &\leqslant \left\|\mathcal{T}'F^{=d}\right\|_2^2 \left\|G^{=d}\right\|_2^2\\
&= \left\|G^{=d}\right\|_2^2 \langle F^{=d}, \mathcal{T}'^{*}\mathcal{T}'F^{=d}\rangle\\
&= \left\|G^{=d}\right\|_2^2 \langle F^{=d}, \Phi F^{=d}\rangle\\
&\leqslant \left(q^{-2d\delta\ell} + 3q^{d-n}\right)\left\|F^{=d}\right\|_2^2\left\|G^{=d}\right\|_2^2\\
&\leqslant \left(q^{-2d\delta\ell} + 3q^{d-n}\right)\alpha\beta,
\end{aligned}$$

where the third transition uses Lemma A.12 and the fact that $F^{=d}$ is basis invariant by Lemma A.10, and finally the fourth transition uses Lemma A.13. Thus, the total contribution from the $d > r$ terms is

$$\sum_{d=r+1}^{2\ell}\left|\langle \mathcal{T}'F^{=d}, G^{=d}\rangle\right| \leqslant \sum_{d=r+1}^{2\ell} 2q^{-d\delta\ell}\sqrt{\alpha\beta} \leqslant q^{-r\delta\ell}\sqrt{\alpha\beta}.$$

Next, we bound the contribution from $d \leqslant r$ by bounding each term separately. Fix a $d \leqslant r$. By Lemma A.15 and Holder's inequality we have

$$\left|\langle \mathcal{T}'F^{=d}, G^{=d}\rangle\right| = \left|\langle \mathcal{T}'F^{=d}, G\rangle\right| \leqslant \left\|\mathcal{T}'F^{=d}\right\|_t \|G\|_{t/(t-1)} \leqslant \beta^{(t-1)/t}\left\|F^{=d}\right\|_t \leqslant q^{500d^2t}\beta^{\frac{t-1}{t}}\varepsilon^{\frac{t-1}{t}},$$

where in the last inequality we are using the fact that $\varepsilon \geqslant \alpha$ and F is (r,ε)-pseudo-random, so by Theorem A.7

$$\left\|F^{=d}\right\|_t \leqslant q^{500d^2t}\varepsilon^{\frac{t-1}{t}}.$$

Altogether, this shows

$$\langle \mathcal{T}'F, G\rangle \leqslant q^{O_{t,r}(1)}\beta^{\frac{t-1}{t}}\varepsilon^{\frac{t-1}{t}} + q^{-r\delta\ell}\sqrt{\alpha\beta}. \qquad \square$$

A.3 Reduction to the Bilinear Scheme

We are now ready to prove Lemma 2.3. As in the statement of Lemma 2.3, let $F \in L_2(\mathsf{Grass}_q(n, 2\ell))$ and $G \in L_2(\mathsf{Grass}_q(n, 2(1-\delta)\ell))$ be Boolean functions, and suppose that F is (r, ε)-pseudo-random. Define the Boolean functions $F' \in L_2\left(\mathbb{F}_q^{n\times 2\ell}\right)$, $G' \in \left(L_2(\mathbb{F}_q^{n\times 2\ell})\right)$ by

$$F'(x_1, \ldots, x_{2\ell}) = \begin{cases} F(\mathrm{span}(x_1, \ldots, x_{2\ell})) & \text{if } \dim(\mathrm{span}(x_1, \ldots, x_{2\ell}) = 2\ell, \\ 0, & \text{otherwise,} \end{cases}$$

and

$$G'(x_1, \ldots, x_{2(1-\delta)\ell}) = \begin{cases} G(\mathrm{span}(x_1, \ldots, x_{2(1-\delta)\ell})) & \text{if } \dim(\mathrm{span}(x_1, \ldots, x_{2(1-\delta)\ell}) = 2(1-\delta)\ell, \\ 0, & \text{otherwise.} \end{cases}$$

We note that F' and G' are basis invariant functions. Next, we will prove that F' is $(r, 2\varepsilon)$ pseudo-random, and towards this end we begin with the following lemma that simplifies the type of zoom-ins and zoom-out combinations we have to consider for F'.

Lemma A.17. *For any $(U, V), (X, Y)$ such that $\dim(U) + \mathrm{codim}(X) < 2\ell$ and $\mathsf{Zoom}[(U, V), (X, Y)]$ is nonempty, there are r' linearly independent columns of V, say $v_1, \ldots, v_{r'} \in \mathbb{F}_q^n$ and a subset of linearly independent rows of X, say $X' \in \mathbb{F}_q^{s'\times n}$, such that $r' \leqslant \dim(U)$, $s' \leqslant \mathrm{codim}(X)$ and*

$$\left\|F'_{(U,V),(X,Y)}\right\|_2^2 \leqslant 2 \cdot \mathop{\mathbb{E}}_{M' \in \mathbb{F}_q^{n\times(2\ell - r')}} [F'\left([v_1, \ldots, v_{r'}, M']\right) \mid X'M' = 0],$$

where $[v_1, \ldots, v_{r'}, M'] \in \mathbb{F}_q^{n\times 2\ell}$ is the matrix whose first r columns are $v_1, \ldots, v_r$, and remaining columns are M'.

Proof. Let $r = \dim(U)$ and $s = \mathrm{codim}(X)$. First note that we can assume that the columns of U and V respectively are both nonzero and linearly independent. Indeed, otherwise say $u_i = 0$, then either $v_i = 0$, in which case the ith columns of U and V can be removed, or $v_i \neq 0$ and $\mathsf{Zoom}[(U, V), (X, Y)]$ is an empty set. Otherwise, if, say, $v_i = 0$, then either $u_i = 0$ and again we can ignored the ith columns, or $u_i \neq 0$ and $\mathsf{Zoom}[(U, V), (X, Y)]$ consists of matrices whose columns are not linearly independent. In this case $F'_{(U,V),(X,Y)}$ is identically 0 and the statement is trivially true. Similarly, if the columns of V are not linearly independent, then $\mathsf{Zoom}[(U, V), (X, Y)]$ consists of matrices whose columns are not linearly independent, and again $F'_{(U,V),(X,Y)}$ is identically 0. Finally, if the columns of U are not linearly independent, then either $\mathsf{Zoom}[(U, V), (X, Y)]$ is empty or there must be some i such that both u_i and v_i are linear combinations of the other columns in U and V respectively, with the same coefficients. In this case, we can remove the ith columns of U and V without changing $\mathsf{Zoom}[(U, V), (X, Y)]$.

Now suppose that the columns of U and V are nonzero and linearly independent, and let $A \in \mathbb{F}_q^{2\ell\times 2\ell}$ be a full rank matrix such that $AU_i = e_i$ for $1 \leqslant i \leqslant r$. Let $Y' \in \mathbb{F}_q^{n\times(2\ell - r)}$ denote the last $2\ell - r$ columns of YA^{-1}, and let Y'' denote the first r columns of YA^{-1}. Since we assumed $\mathsf{Zoom}[(U, V), (X, Y)]$ is nonempty, we must have $X[v_1,\dots,v_r]=Y''$. Then by the fact that F' is basis invariant we get that

$$\begin{aligned}\left\|F'_{(U,V),(X,Y)}\right\|_2^2 &= \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}}[F'(M)\mid MU=V, XM=Y]\\ &= \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}}[F'(MA^{-1})\mid MU=V, XM=Y]\\ &= \mathop{\mathbb{E}}_{M\in\mathbb{F}_q^{n\times 2\ell}}[F'(M)\mid MAU=V, XM=Y']\\ &= \mathop{\mathbb{E}}_{M'\in\mathbb{F}_q^{n\times(2\ell-r')}}[F'([v_1,\dots,v_r,M'])\mid XM'=Y'].\end{aligned}$$

To complete the proof, we show to reduce to the case that Y' is the zero matrix. First note that, using the same reasoning as we did for U and V, we can assume that the nonzero rows Y' are linearly independent and the rows of X are linearly independent

Suppose that $y'_1,\dots,y'_a\in\mathbb{F}_q^{2\ell-r}$ are the nonzero (and linearly independent) rows of Y', while the remaining rows are $y'_{a+1},\dots,y'_s=0$. Let $Y''\in\mathbb{F}_q^{a\times(2\ell-r)}$ be the first a rows of Y', which are nonzero, let $X'\in\mathbb{F}_q^{a\times n}$ denote the first a rows of X, and let $X''\in\mathbb{F}_q^{(2\ell-a)\times n}$ denote rows $a+1$ through s of X. For any $Z=(z_1,\dots,z_a)\in\mathbb{F}_q^{a\times(2\ell-r)}$ with a linearly independent rows let $A_Z\in\mathbb{F}_q^{(2\ell-r)\times(2\ell-r)}$ be the full rank matrix such that $Y''A_Z=Z$. Then, for any linearly independent $z_1,\dots,z_a\in\mathbb{F}_q^{2\ell-r}$,

$$\begin{aligned}&\mathop{\mathbb{E}}_{M'\in\mathbb{F}_q^{n\times(2\ell-r')}}[F'([v_1,\dots,v_r,M'])\mid XM'=Y']\\ &\qquad= \mathop{\mathbb{E}}_{M'}[F'([v_1,\dots,v_r,M'])\mid X'M'=Y'', X''M'=0]\\ &\qquad= \mathop{\mathbb{E}}_{M'}[F'([v_1,\dots,v_r,M'A_Z^{-1}])\mid X'M'=Y'', X''M'=0]\\ &\qquad= \mathop{\mathbb{E}}_{M'}[F'([v_1,\dots,v_r,M'])\mid X'M'=Z, X''M'=0].\end{aligned}$$

In the second transition, we used the fact that F' is basis invariant and multiplied its input by the matrix whose top left $r\times r$ minor is the identity, its bottom right $(2\ell-r)\times(2\ell-r)$ is A_Z^{-1}, and the rest of the entries are 0. Since the above holds for any Z with a-linearly independent rows, letting E denote the event that $X'M'$ has a linearly independent rows, it follows that

$$\mathop{\mathbb{E}}_{M'\in\mathbb{F}_q^{n\times(2\ell-r')}}[F'([v_1,\dots,v_r,M'])\mid XM'=Y'] = \mathop{\mathbb{E}}_{M'}[F'([v_1,\dots,v_r,M'])\mid E\wedge X''M'=0],$$

and

$$\begin{aligned}\mathop{\mathbb{E}}_{M'\in\mathbb{F}_q^{n\times(2\ell-r')}}[F'([v_1,\dots,v_r,M'])\mid X''M'=0] &\geqslant \Pr_{M'\in\mathbb{F}_q^{n\times(2\ell-r')}}[E\mid X''M'=0]\\ &\qquad\cdot\mathop{\mathbb{E}}_{M'}[F'([v_1,\dots,v_r,M'])\mid E\wedge XM'=Y'].\end{aligned}$$

Finally since $\Pr[E\mid X''M'=0]\geqslant\frac{1}{2}$ (as it is the probability of choosing $a<2\ell-r$ linearly independent vectors in $\mathbb{F}_q^{2\ell-r}$), we have,

$$\begin{aligned}\left\|F'_{(U,V),(X,Y)}\right\|_2^2 &= \mathop{\mathbb{E}}_{M'\in\mathbb{F}_q^{n\times(2\ell-r')}}[F'([v_1,\dots,v_r,M'])\mid E\wedge XM'=Y']\\ &\leqslant 2\mathop{\mathbb{E}}_{M'}[F'([v_1,\dots,v_r,M'])\mid X''M'=0],\end{aligned}$$

and the proof is concluded. $\square$

As an immediate consequence of Lemma A.17

Lemma A.18. *If F is (r, ε)-pseudo-random then F' is $(r, 2\varepsilon)$-pseudo-random.*

Proof. Fix any (U, V) and (X, Y) such that $\dim(U) + \operatorname{codim}(X) = r$. Using Lemma A.17, there are linearly independent $v_1, \ldots, v_{r'} \in \mathbb{F}_q^n$ and $X' \in \mathbb{F}_q^{s' \times n}$ with linearly independent rows such that

$$\begin{aligned}
&\left\|F'_{(U,V),(X,Y)}\right\|_2^2 \\
&\qquad \leqslant 2 \cdot \mathop{\mathbb{E}}_{M' \in \mathbb{F}_q^{n \times (2\ell - r')}} \left[F'\left([v_1, \ldots, v_{r'}, M']\right) \mid X'M' = 0\right] \\
&\qquad \leqslant 2 \cdot \mathop{\mathbb{E}}_{M' \in \mathbb{F}_q^{n \times (2\ell - r')}} \left[F'\left([v_1, \ldots, v_{r'}, M']\right) \mid X'M' = 0, \dim(\operatorname{im}([v_1, \ldots, v_r, M'])) = 2\ell\right] \\
&\qquad = 2 \cdot \mathop{\mathbb{E}}_{M' \in \mathbb{F}_q^{n \times (2\ell - r')}} \left[F\left(\operatorname{im}([v_1, \ldots, v_{r'}, M'])\right) \mid X'M' = 0, \dim(\operatorname{im}([v_1, \ldots, v_r, M'])) = 2\ell\right],
\end{aligned}$$

where in the second transition we are using the fact that $F'(M) = 0$ for all M such that $\dim(\operatorname{im}(M)) < 2\ell$, and in the third transition we are using the definition of F'. We will bound the final term by using the pseudo-randomness of F.

Choosing $M' \in \mathbb{F}_q^{n \times (2\ell - r')}$ uniformly conditioned on $X'M' = 0$, and $\dim(\operatorname{im}([v_1, \ldots, v_r, M']))$, we claim that $\operatorname{im}([v_1, \ldots, v_r, M'])$ is a uniformly random 2ℓ-dimensional subspace in $\mathsf{Zoom}[Q, Q \oplus H]$, where H is the codimension s subspace that is dual to the rows of X'. To see why, first note that it is clear $\operatorname{im}([v_1, \ldots, v_r, M']) \in \mathsf{Zoom}[Q, Q \oplus H]$. Additionally, each $L \in \mathsf{Zoom}[Q, Q \oplus H]$ has an equal number of $M' \in \mathbb{F}_q^{n \times (2\ell - r')}$ such that

$$L = \operatorname{im}([v_1, \ldots, v_r, M']),$$

and therefore has an equal chance of being selected. It follows that choosing $M' \in \mathbb{F}_q^{n \times (2\ell - r')}$ uniformly conditioned on $X'M' = 0$, and $\dim(\operatorname{im}([v_1, \ldots, v_r, M']))$, $\operatorname{im}([v_1, \ldots, v_r, M'])$ is a uniformly random 2ℓ-dimensional subspace in $\mathsf{Zoom}[Q, Q \oplus H]$. As a result,

$$\begin{aligned}
&\left\|F'_{(U,V),(X,Y)}\right\|_2^2 \\
&\qquad \leqslant 2 \cdot \mathop{\mathbb{E}}_{M' \in \mathbb{F}_q^{n \times (2\ell - r')}} \left[F\left(\operatorname{im}([v_1, \ldots, v_{r'}, M'])\right) \mid X'M' = 0, \dim(\operatorname{im}([v_1, \ldots, v_r, M'])) = 2\ell\right] \\
&\qquad = 2 \cdot \mathop{\mathbb{E}}_{L \in \mathsf{Zoom}[Q, Q \oplus H]} [F(L)] \\
&\qquad \leqslant 2\varepsilon,
\end{aligned}$$

where in the last transition we use the fact that F is (r, ε)-pseudo-random and $\dim(Q) + \operatorname{codim}(Q \oplus H) \leqslant r' + s \leqslant r$. $\square$

Next, we note that the values of $\langle \mathcal{T}F, G\rangle$ and $\langle \mathcal{T}'F, G\rangle$ are similar.

Lemma A.19. *We have*

$$\langle \mathcal{T}F, G\rangle \leqslant 2\langle \mathcal{T}'F', G'\rangle.$$

Proof. We have,

$$\begin{aligned}
\langle \mathcal{T}' F', G' \rangle &= \mathop{\mathbb{E}}_{M \in \mathbb{F}_q^{n\times 2\ell}, A \in \mathbb{F}_q^{2\ell \times 2(1-\delta)\ell}} [F'(M) \cdot G'(MA) \mid \operatorname{rank}(A) = 2(1-\delta)\ell] \\
&\geqslant \Pr_{M,A}[\operatorname{rank}(M) = 2\ell, \operatorname{rank}(MA) = 2(1-\delta)\ell \mid \operatorname{rank}(A) = 2(1-\delta)\ell] \\
&\cdot \mathop{\mathbb{E}}_{M,A}[F'(M) \cdot G'(MA) \mid \operatorname{rank}(M) = 2\ell, \operatorname{rank}(MA) = 2(1-\delta)\ell, \operatorname{rank}(A) = 2(1-\delta)\ell] \\
&= \frac{1}{2} \mathop{\mathbb{E}}_{M,A}[F'(M) \cdot G'(MA) \mid \operatorname{rank}(M) = 2\ell, \operatorname{rank}(MA) = 2(1-\delta)\ell, \operatorname{rank}(A) = 2(1-\delta)\ell] \\
&= \frac{1}{2} \mathop{\mathbb{E}}_{M,A}[F(\operatorname{im}(M)) \cdot G(\operatorname{im}(MA)) \mid \operatorname{rank}(M) = 2\ell, \operatorname{rank}(MA) = 2(1-\delta)\ell, \operatorname{rank}(A) = 2(1-\delta)\ell].
\end{aligned}$$

To finish the proof, notice that in the conditional distribution $(\operatorname{im}(M), \operatorname{im}(MA))$ in the last term, $\operatorname{im}(M)$ is a uniform $L \in \mathsf{Grass}_q(n, 2\ell)$ and $\operatorname{im}(MA)$ is a uniform $L' \in \mathsf{Grass}_q(n, 2(1-\delta)\ell)$ such that $L' \subseteq L$. Therefore,

$$\langle \mathcal{T}' F', G' \rangle \geqslant \frac{1}{2} \mathop{\mathbb{E}}_{L,L'}[F(L) \cdot G(L') \mid L \supseteq L'] = \frac{1}{2} \langle \mathcal{T} F, G \rangle. \qquad \square$$

Lemma 2.3 follows by combining Lemma A.18 and Lemma A.19.

Proof of Lemma 2.3. Suppose $F \in L_2(\mathsf{Grass}_q(n, 2\ell)), G \in L_2(\mathsf{Grass}_q(n, 2(1-\delta)\ell))$ have expectations α, β respectively, and suppose that F is (r, ε) pseudo-random. Define the associated functions $F' \in L_2\left(\mathbb{F}_q^{n \times 2\ell}\right)$ and $G' \in L_2\left(\mathbb{F}_q^{n \times 2(1-\delta)\ell}\right)$ as above. It is clear that,

$$\left\|F'\right\|_2^2 \leqslant \|F\|_2^2 = \alpha, \quad \left\|G'\right\|_2^2 \leqslant \|G\|_2^2 = \beta.$$

Furthermore, by Lemma A.19, we have that

$$\langle \mathcal{T} F, G \rangle \leqslant 2 \langle \mathcal{T} F', G' \rangle.$$

By Lemma A.18 F' is $(r, 2\varepsilon)$-pseudo-random, and applying Lemma A.16 we get that

$$\langle \mathcal{T} F, G \rangle \leqslant 2 \langle \mathcal{T}' F', G' \rangle \leqslant q^{O_{t,r}(1)} \beta^{(t-1)/t} \varepsilon^{2t/(2t-1)} + q^{-r\delta\ell} \sqrt{\alpha\beta}. \qquad \square$$

A.4 Proof of Lemma 2.4

We will show that if a set of ℓ'-dimensional subspaces $\mathcal{L}^\star \subseteq \mathsf{Grass}_q(V^\star, \ell')$ is pseudo-random, then it must "evenly cover" the space V in the sense that there are very few points $z \in V$ such that $\mu_z(\mathcal{L}^\star)$ significantly deviates from $\mu(\mathcal{L}^\star)$. We will require the following result from [EKL23b].

Theorem A.20. *[EKL23b, Theorem 1.12] If $F' \in L_2\left(\mathbb{F}_q^{n \times \ell'}\right)$ is a Boolean function which is $(1, \varepsilon)$-global, then for all powers of $t \geqslant 4$ that are powers of 2 it holds that*

$$\left\|F'^{=1}\right\|_2^2 \leqslant q^{460t} \left\|F\right\|_2^{2-\frac{2}{t}} \varepsilon.$$

Lemma A.21. *Let $\mathcal{L}^\star \subseteq \mathsf{Grass}_q(V^\star, \ell')$ have $\mu(\mathcal{L}^\star) = \eta \geqslant q^{-C\ell'}$ for some large constant C, and set $Z = \{z \in V^\star \mid |\mu_z(\mathcal{L}^\star) - \eta| \leqslant \frac{\eta}{10}\}$. If $\mathcal{L}^\star$ is $(1, q^{c\ell'}\eta)$-pseudo-random for some $0 < c < 1$, then*

$$|Z| \geqslant \left(1 - q^{\frac{\ell'}{2}}\right) |V^\star|.$$

Proof. Let $\dim(V^\star) = n$, let $F \in L_2(\mathsf{Grass}_q(V^\star, \ell'))$ be the indicator function for $\mathcal{L}^\star$, and let $F' \in L_2\left(\mathbb{F}_q^{n\times \ell'}\right)$ be the associated function given by

$$F'(x_1, \ldots, x_{\ell'}) = \begin{cases} F(\mathrm{span}(x_1, \ldots, x_{\ell'})) & \text{if } \dim(\mathrm{span}(x_1, \ldots, x_{2\ell})) = \ell', \\ 0 & \text{otherwise.} \end{cases}$$

By Lemma A.18, F' is $(1, 2q^{c\ell'}\eta)$-pseudo-random. For any point $x \in V^\star$, we have

$$\mu_z(\mathcal{L}^\star) = \mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}\in V^\star}\left[F'(x_1, \ldots, x_{\ell'-1}, z) \mid \dim(\mathrm{span}(x_1, \ldots, x_{\ell'-1}, z)) = \ell'\right],$$

so it follows that

$$\left|\mu_z(\mathcal{L}^\star) - \mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}}\left[F'(x_1, \ldots, x_{\ell'-1}, z)\right]\right| \leqslant \frac{q^{\ell'}}{q^n} \quad \text{and} \quad \left|\mu(\mathcal{L}^\star) - \left\|F'\right\|_2^2\right| \leqslant \frac{q^{\ell'}}{q^n}.$$

Thus

$$\left|\mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}}\left[F'(x_1, \ldots, x_{\ell'-1}, z)\right] - \left\|F'\right\|_2^2\right| \geqslant |\mu_z(\mathcal{L}^\star) - \mu(\mathcal{L}^\star)| - \frac{q^{\ell'}}{q^n}. \tag{19}$$

We will now relate this quantity to the level one weight of F' and apply Lemma A.7 to bound the level one weight of F'. Note that

$$\begin{aligned}
\mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}}\left[F'(x_1, \ldots, x_{\ell'-1}, z)\right] &= \sum_{S=(s_1,\ldots,s_{\ell'})\in V^\star} \widehat{F'}(S) \mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}}\left[\chi_S(x_1, \ldots, x_{\ell'-1}, z)\right] \\
&= \sum_{S=(s_1,\ldots,s_{\ell'})\in V^\star} \widehat{F'}(S)\chi_{s_{\ell'}}(z) \prod_{i=1}^{\ell'-1} \mathop{\mathbb{E}}_{x_i}\left[\chi_{s_i}(x_i)\right].
\end{aligned}$$

Now note that $\mathbb{E}_{x_i}\left[\chi_{s_i}(x_i)\right] = 0$ if s_i is not the zero vector, and $\mathbb{E}_{x_i}\left[\chi_{s_i}(x_i)\right] = 1$ if s_i is the zero vector. Thus,

$$\mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}}\left[F'(x_1, \ldots, x_{\ell'-1}, z)\right] = \widehat{F'}(0, \ldots, 0) + \sum_{a\in V^\star, a\neq 0} \widehat{F}(0, \ldots, 0, a)\chi_a(z),$$

and using the fact that $\widehat{F'}(0, \ldots, 0) = \mathbb{E}[F'] = \|F'\|_2^2$,

$$\left(\mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}}\left[F'(x_1, \ldots, x_{\ell'-1}, z)\right] - \left\|F'\right\|_2^2\right)^2 = \left(\sum_{a\in V^\star, a\neq 0} \widehat{F}(0, \ldots, 0, a)\chi_a(z)\right)^2.$$

Therefore, we get by (19) that

$$\begin{aligned}
\mathop{\mathbb{E}}_{z\in V^\star}\left[(\mu_z(\mathcal{L}^\star) - \eta)^2\right] &\leqslant \mathop{\mathbb{E}}_{z\in V^\star}\left[\left|\mathop{\mathbb{E}}_{x_1,\ldots,x_{\ell'-1}}\left[F'(x_1, \ldots, x_{\ell'-1}, z)\right] - \left\|F'\right\|_2^2\right|^2\right] + 5\frac{q^{\ell'}}{q^n} \\
&\leqslant \mathop{\mathbb{E}}_{z}\left[\left|\sum_{a\in V^\star, a\neq 0} \widehat{F}(0, \ldots, 0, a)\chi_a(z)\right|^2\right] + 5\frac{q^{\ell'}}{q^n} \\
&= \sum_{a\in V^\star, a\neq 0} |\widehat{F}(0, \ldots, 0, a)|^2 + 5\frac{q^{\ell'}}{q^n}.
\end{aligned}$$

Next since F' is basis invariant, using Lemma A.8, we have that for all $\alpha_1, \ldots, \alpha_{\ell'} \in \mathbb{F}_q$ that are not all zero,

$$\widehat{F'}\left(\alpha_1 a, \ldots, \alpha_{\ell'} a\right) = \widehat{F'}\left(0, \ldots, 0, a\right).$$

It follows that

$$\sum_{a \in V^\star, a \neq 0} |\widehat{F'}(0, \ldots, 0, a)|^2 = \frac{1}{q^{\ell'} - 1} \sum_{\mathrm{rank}(S)=1} |\widehat{F'}(S)|^2 = \frac{\left\|F'^{=1}\right\|_2^2}{q^{\ell'} - 1}.$$

Using Theorem A.20 with $d = 1$, along the fact that F' is $(1, q^{c\ell'}\eta)$-pseudo-random, we get

$$\mathop{\mathbb{E}}_{z \in V^\star}[(\mu_z(\mathcal{L}^\star) - \eta)^2] \leqslant \frac{\left\|F'^{=1}\right\|_2^2}{q^{\ell'} - 1} + 5\frac{q^{\ell'}}{q^n} \leqslant \frac{q^{460t} q^{c\ell'} \eta^{2-\frac{2}{t}}}{q^{\ell'} - 1},$$

for any $t \geqslant 4$ that is a power of 2. By Markov's inequality it follows that

$$\frac{|\overline{Z}|}{q^n} \cdot \frac{\eta^2}{100} \leqslant \frac{100 q^{460t} q^{c\ell'}}{\eta^{\frac{2}{t}} q^{\ell'}} \eta^2 \leqslant \frac{100 q^{460t} q^{c\ell'}}{q^{\left(1-c-\frac{2C}{t}\right)\ell'}} \eta^2 \leqslant \frac{q^{-\frac{\ell'}{2}}}{100} \eta^2,$$

where we take t to be a power of 2 large enough so that $1 - c + \frac{2C}{t} \geqslant \frac{1}{2}$. Dividing by η^2 finishes the proof. $\square$

B Proof of Theorem 5.2

In order to prove Theorem 5.2 we will find the subspaces, Q, one at at time by using Theorem 5.1. We let $\mathcal{Q}$ denote the set of all Q's collected thus far. Each time a new subspace Q is added to $\mathcal{Q}$, we randomize the assignment $T_1[L]$ for all 2ℓ-dimensional $L \supset Q$. At a high level, the effect of this randomization is that there is only a little agreement between any linear function and the assignments on subspaces containing Q, thus these entries are essentially "deleted".

More formally, we construct the set $\mathcal{Q}$ of subspaces as follows. Initially set $\tilde{T}_1 = T_1$, $\mathcal{Q} = \emptyset$, and $X = \emptyset$. Recall that initially $\tilde{T}_1$ and T_2 are ε-consistent for $\varepsilon \geqslant 2q^{-2\ell(1-1000\delta)}$. While $\tilde{T}_1$ and T_2 are at least $\varepsilon/2$-consistent, do the following.

1. Let $Q \subset W$ be subspaces guaranteed by Theorem 5.1. That is, Q and W satisfy $\dim(Q) + \mathrm{codim}(W) = r$ and there exists linear $g_{Q,W} :\to W \to \mathbb{F}_q$ such that

$$\Pr_{L \in \mathsf{Grass}(n, 2\ell)}[g_{Q,W}|_L = \tilde{T}_1[L] \mid Q \subseteq L \subseteq W] \geqslant \varepsilon'.$$

2. Set $\mathcal{Q} \leftarrow \mathcal{Q} \cup \{Q\}$.
3. Set $X \leftarrow X \cup \{L \mid Q \subseteq L \subseteq W\}$.
4. For each $L \in X$ independently, choose $\tilde{T}_1[L]$ uniformly among all linear functions on L.

We have the following claim regarding the re-assignment phase.

Claim B.1. *With probability $1 - e^{-\Omega(q^{\ell n})}$ over the random assignment on X, for every Q, W such that $Q \subset W$ and $\dim(Q) + \mathrm{codim}(W) = r$ at least one of the following holds:*

1. *Less than $q^{-2\ell}$-fraction of $L \in \mathsf{Zoom}[Q, W]$ are in X,*
2. *For every linear function $g_{Q,W} : W \to \mathbb{F}_q$,*

$$\Pr_{L \in X, L \in \mathsf{Zoom}[Q,W]}[g_{Q,W} \equiv \tilde{T}_1[L]] \leqslant q^{1-2\ell}$$

Proof. Note that the first item has nothing to do with the random assignment over X, so we need only show that if the first item is false then the second item must be true. Suppose that the first item does not hold and that X contains at least $q^{-2\ell}$-fraction of $L \in \mathsf{Zoom}[Q, W]$.

Fix $Q, W, g_{Q,W}$ with the parameters above and let $A = \mathsf{Zoom}[Q, W] \cap X$, and suppose that the first item does not hold. We will show that in this case, the second item holds.

For each $L \in A$, let Z_L denote the indicator variable that takes value 1 if $g_{Q,W}|_L = \tilde{T}_1[L]$ and 0 otherwise. The expectation of Z_L is $q^{-2\ell}$, and by a Chernoff bound, the desired probability is bounded by

$$\Pr_{L \in A}\left[\frac{1}{|A|}\sum_{L \in A} Z_L \geqslant q^{1-2\ell}\right] \leqslant e^{-q^{-2\ell+2}|A|/6}.$$

By assumption, $|A| \geqslant q^{-2\ell}\left|\{L \in \mathsf{Zoom}[Q, W], \dim(L) = 2\ell\}\right| \geqslant q^{-2\ell}q^{(2\ell-r)(n-r-2\ell)}$. Thus, using a union bound over all $Q, W, g_{Q,W}$, the probability that there exist a bad triple is at most,

$$(r+1)q^{nr}q^{n}e^{-q^{-4\ell+2}q^{(2\ell-r)(n-r-2\ell)}/6} \leqslant e^{-\Omega(q^{\ell n})}.$$ □

We now analyze the process. Note that using Chernoff's bound, with probability $1 - e^{-\Omega(q^{\ell n})}$ over the randomization step, the probability $\Pr_{L \in \mathsf{Grass}(n,2\ell)}[g_{Q,W}|_L = \tilde{T}_1[L] \mid Q \subseteq L \subseteq W]$ drops from at least ε' to at most $q^{1-2\ell}$. In that case, the measure of X increases by at least

$$(\varepsilon' - q^{1-2\ell})q^{-rn} \geqslant q^{-O(rn)}.$$

Doing a union bound over the steps, it follows that with probability $1 - e^{-\Omega(q^{\ell n})}q^{O(rn)} = 1 - o(1)$ the process terminates within $q^{O(rn)}$ steps.

Note that it is possible that the same subspace Q is added multiple times (with different zoom-outs) in the process above, so we clarify that $\mathcal{Q}$ is considered as a set without repeats. Also note that with probability $1 - o(1)$, for each $Q \in \mathcal{Q}$, W and $g_{Q,W}$ found in the process it holds that

$$\Pr_{L \in \mathsf{Grass}(n,2\ell)}[g_{Q,W}|_L = T_1[L] \mid Q \subseteq L \subseteq W] \geqslant \frac{1}{2}\varepsilon' \tag{20}$$

(the point being is that the agreement now is compared to the original T_1 and not to $\tilde{T}_1$). Indeed, considering the step Q, W and $g_{Q,W}$ were found, $g_{Q,W}$ had agreement at least ε' with $\tilde{T}_1$ on $\mathsf{Zoom}[Q, W]$ at that point, and by Claim B.1 with probability $1 - e^{-\Omega(q^{\ell n})}$ at most $q^{1-2\ell} \leqslant \varepsilon'/2$ of that agreement came from $L \in X$. Thus, by union bound over all of the steps, with probability $1 - q^{O(rn)}q^{-\Omega(\ell n)} = 1 - o(1)$ it follows that (20) holds for every Q, W and $g_{Q,W}$ found throughout the process.

The following claim shows that at the end of the process the number of Q's found in the process is large, thereby finishing the proof of Theorem 5.2.

Claim B.2. *There exists some $0 \leqslant r_1 \leqslant r$ such that $\mathcal{Q}$ contains at least a $q^{-5\ell^2}$-fraction of all r_1-dimensional subspaces.*

Proof. At the end of the process, the consistency has dropped by at least $\varepsilon'/2$, so the probability over edges (L, R) that L was reassigned must be at least $\varepsilon'/2$. For each $0 \leqslant r_1 \leqslant r$, let N_{r_1} be the number of Q of dimension r_1 in $\mathcal{Q}$.

For each Q of dimension r_1, the fraction of 2ℓ-dimensional L's that are reassigned due to Q being added to $\mathcal{Q}$ is at most the fraction of 2ℓ-dimensional subspaces that contain Q. This is,

$$\frac{\begin{bmatrix} n \\ 2\ell - r_1 \end{bmatrix}_q}{\begin{bmatrix} n \\ 2\ell \end{bmatrix}_q} \leqslant \frac{q^{n(2\ell - r_1)}}{q^{2\ell(n-2\ell)}} = q^{4\ell^2 - r_1 n}.$$

It follows that there must be some r_1 such that

$$N_{r_1} q^{4\ell^2 - r_1 n} \geqslant \frac{\varepsilon}{r+1}.$$

Rearranging this inequality, we get that

$$N_{r_1} \geqslant \frac{\varepsilon}{(r+1)q^{4\ell^2}} q^{r_1} n \geqslant q^{-5\ell^2} \begin{bmatrix} n \\ r_1 \end{bmatrix}_q.$$

Thus there exists an r_1 such that $\mathcal{Q}$ contains at least a $q^{-5\ell^2}$-fraction of all r_1-dimensional subspaces. $\square$

C The Covering Property

Fix a question U to the first prover. Recall that we set

$$k = q^{2(1+c)\ell} \quad \text{and} \quad \beta = q^{-2(1+2c/3)\ell},$$

where $0 < c < 1$ is some small constant close to 0 and set $\eta = q^{-100\ell^{100}}$, and recall that the distributions $\mathcal{D}$ and $\mathcal{D}'$ are defined as follows.

$\mathcal{D}$:

- Choose $x_1, \ldots, x_{2\ell} \in \mathbb{F}_q^U$ uniformly.
- Output the list $(x_1, \ldots, x_{2\ell})$.

$\mathcal{D}'$:

- Choose $V \subseteq U$ according to the Outer PCP.
- Choose $x'_1, \ldots, x'_{2\ell} \in \mathbb{F}_q^V$ uniformly, and lift these vectors to $\mathbb{F}_q^U$ by inserting 0's into the missing coordinates.
- Choose $w_1, \ldots, w_{2\ell} \in H_U$ uniformly, and set $x_i = x'_i + w_i$ for $1 \leqslant i \leqslant 2\ell$.
- Output the list $(x_1, \ldots, x_{2\ell})$.

We also restate Lemma 5.4 below as a reminder.

Lemma 5.4. *Let η be a parameter such that $q^{-100\ell^{100}} \leqslant \eta \leqslant 1/2$. There exists a small set $E \subseteq \left(\mathbb{F}_q^U\right)^{2\ell}$ such that both $\mathcal{D}(E)$ and $\mathcal{D}'(E)$ are at most η^{40}, and for all $(x_1, \ldots, x_{2\ell}) \notin E$ we have*

$$0.9 \leqslant \frac{D(x_1, \ldots, x_{2\ell})}{D'(x_1, \ldots, x_{2\ell})} \leqslant 1.1.$$

C.1 Proof of Lemma 5.4

For $x_1, \ldots, x_{2\ell} \in \mathbb{F}_q^{|U|}$, let us view $x_1, \ldots, x_{2\ell}$ as the rows of a $2\ell \times 3k$ matrix, and split the columns of this matrix into k blocks - each consisting of 3 consecutive columns. Then let $s(x_1, \ldots, x_{2\ell})$ be the number of blocks where at least two of the columns are equal, and set $p = 3q^{-2\ell} - 2q^{-4\ell}$. The idea is that $s(x_1, \ldots, x_{2\ell})$ should correspond to the number of equations where we drop variables in the Outer PCP, while p is the probability that a fixed block has at least two columns equal to each other. Also let $s'(x_1, \ldots, x_{2\ell})$ be the number of blocks where all 3 columns are equal, and let $p' = q^{-4\ell}$, be the probability that a fixed block has all three columns equal.

We define the set E as follows. Set

$$E_1 = \left\{(x_1, \ldots, x_{2\ell}) \in \left(\mathbb{F}_q^U\right)^{2\ell} \mid |s(x_1, \ldots, x_{2\ell}) - pk| > 50\sqrt{pk\log(1/\eta)}\right\}, \tag{21}$$

$$E_2 = \left\{(x_1, \ldots, x_{2\ell}) \in \left(\mathbb{F}_q^U\right)^{2\ell} \mid s'(x_1, \ldots, x_{2\ell}) > \ell^{100}\right\}, \tag{22}$$

and $E = E_1 \cup E_2$. By a Chernoff bound,

$$\mathcal{D}(E_1) = \Pr_{\mathcal{D}}\left[|s(x_1, \ldots, x_{2\ell}) - pk| > 50\sqrt{pk\log(1/\eta)}\right] \leqslant \eta^{50},$$

where recall that $\eta = q^{-100\ell^{100}}$. Also, by our setting of βk, we have $\beta k = q^{2c\ell/3}$, while $pk = O(q^{2c\ell})$, so the same Chernoff bound holds for $\mathcal{D}'$,

$$\mathcal{D}'(E_1) = \Pr_{\mathcal{D}'}\left[|s(x_1, \ldots, x_{2\ell}) - pk| > 50\sqrt{pk\log(1/\eta)}\right] \leqslant \eta^{50},$$

so $\mathcal{D}'(E_1) \leqslant \eta^{45}$. Indeed, the actual expectation of $s(x_1, \ldots, x_{2\ell})$ under $\mathcal{D}'$ is $(1-\beta)pk + \beta k$ and this differs from pk by only $O(\beta k) = O(pk)$.

For the measure of E_2 we have,

$$\mathcal{D}(E_2) = \Pr_{D}[s'(x_1, \ldots, x_{2\ell}) > \ell^{100}] \leqslant \binom{k}{\ell^{100}} p'^{\ell^{100}} \leqslant (kp')^{\ell^{100}} \leqslant \eta^{100},$$

where in the middle term, the first factor is the number of ways to choose ℓ^{100} blocks and the second factor is the probability that all of these blocks have all three columns equal. Similarly,

$$\mathcal{D}'(E_2) = \Pr_{D'}[s'(x_1, \ldots, x_{2\ell}) > \ell^{100}] \leqslant \binom{k}{\ell^{100}} \left((1-\beta)p' + \beta\frac{1}{q^{2\ell}}\right)^{\ell^{100}} \leqslant k^{\ell^{100}} \left(q^{-4\ell}\right)^{\ell^{100}} \leqslant \eta^{100}.$$

Putting everything together, we get that

$$\mathcal{D}(E) \leqslant \mathcal{D}(E_1) + \mathcal{D}(E_2) \leqslant \eta^{40} \quad \text{and} \quad \mathcal{D}'(E) \leqslant \mathcal{D}'(E_1) + \mathcal{D}'(E_2) \leqslant \eta^{40}. \tag{23}$$

We next show that the probability measures $\mathcal{D}$ and $\mathcal{D}'$ assign roughly the same measure to each $x \notin E$. Fix $(x_1, \ldots, x_{2\ell}) \notin E$. It is clear that $D(x_1, \ldots, x_{2\ell}) = q^{-(2\ell)3k}$, where we use $|U| = 3k$. Let $s = s(x_1, \ldots, x_{2\ell})$ and $s' = s'(x_1, \ldots, x_{2\ell})$. Then,

$$\begin{aligned} D'(x_1, \ldots, x_{2\ell}) &= ((1-\beta)q^{-3\cdot 2\ell})^{k-s-s'} \left((1-\beta)q^{-3\cdot 2\ell} + \frac{\beta}{3}q^{-2\cdot\ell}\right)^s \left((1-\beta)q^{-3\cdot 2\ell} + \beta q^{-2\cdot 2\ell}\right)^{s'} \\ &= q^{-2\ell\cdot 3k}(1-\beta)^{k-s-s'} \left(1-\beta+\frac{\beta}{3}q^{2\ell}\right)^s (1-\beta+\beta q^{2\ell})^{s'} \end{aligned} \tag{24}$$

In the first equality, the first term is the probability of choosing the blocks that have three all distinct columns. Then, $(1-\beta)$ is the probability that no variables are dropped, and $q^{-3(2\ell)}$ is the probability of choosing those three particular x_i's in that block. The second term is the probability of choosing the blocks that have exactly two equal columns. Then, $(1-\beta)q^{-3(2\ell)}$ is the probability of having no variables dropped and choosing the three x_i's, and $\frac{\beta}{3}q^{-2(2\ell)}$ is the probability of first having the variable dropped in the column that is not equal to the other two, and then choosing the correct values for the remaining two column values.

We will first show

$$\frac{D'(x_1, \ldots, x_{2\ell})}{D(x_1, \ldots, x_{2\ell})} \geqslant \frac{1}{1.1}.$$

Using Equation (24),

$$\begin{aligned} \frac{D'(x_1, \ldots, x_{2\ell})}{D(x_1, \ldots, x_{2\ell})} &= (1-\beta)^{k-s-s'} \left(1-\beta+\frac{\beta}{3}q^{2\ell}\right)^s (1-\beta+\beta q^{2\ell})^{s'} \\ &\geqslant (1-\beta)^{k-s} \left(1-\beta+\frac{\beta}{3}q^{2\ell}\right)^s \\ &= (1-\beta)^{k-s} \left(1+\beta\left(\frac{q^{2\ell}}{3}-1\right)\right)^s \\ &\geqslant \exp\left(-\beta(k-s) - \beta^2(k-s) + \left(\frac{q^{2\ell}}{3}-1\right)\beta s - \left(\frac{q^{2\ell}}{3}-1\right)^2 \beta^2 s\right) \\ &= \exp\left(-\beta k - \beta^2(k-s) + \frac{q^{2\ell}}{3}\beta s - \left(\frac{q^{2\ell}}{3}-1\right)^2 \beta^2 s\right) \\ &\geqslant \exp\left(-\beta k + \frac{q^{2\ell}}{3}\beta s - \left(\frac{q^{2\ell}}{3}-1\right)^2 \beta^2 s - q^{-2\ell}\right), \end{aligned}$$

where in the fourth transition we use the bound $1+z \geqslant e^{z-z^2}$, which holds for all real numbers z such that $|z|$ is sufficiently small. For our uses, $z = \beta$ and $z = \beta\left(\frac{q^{2\ell}}{3}-1\right)$ are $q^{-2(1+2c/3)\ell}$ and $O(q^{-2c\ell/3})$, and both are sufficiently small. In the last transition we use the fact that $-\beta^2(k-s) \geqslant -\beta^2 k \geqslant -q^{-2\ell}$.

Now write $s = pk - v$ and let us analyze the first two terms in the last line. Plugging this in and using the definition of p we get that

$$-\beta k + \frac{q^{2\ell}}{3}\beta s = -\beta k + \frac{q^{2\ell}}{3}\beta(pk-v) = -\beta k\left(1-\frac{q^{2\ell}}{3}p\right) - \frac{q^{2\ell}}{3}\beta v = -\frac{2}{3}\beta k q^{-2\ell} - \frac{q^{2\ell}}{3}\beta v. \tag{25}$$

Plugging this back into the above,

$$\frac{D'(x_1,\ldots,x_{2\ell})}{D(x_1,\ldots,x_{2\ell})} \geqslant \exp\left(-\left(\frac{q^{2\ell}}{3}-1\right)^2 \beta^2 pk\right)$$
$$\cdot \exp\left(-\frac{q^{2\ell}}{3}\beta v + \left(\frac{q^{2\ell}}{3}-1\right)^2 \beta^2 v\right) \cdot \exp\left(-\frac{2}{3}\beta k q^{-2\ell}\right).$$

Plugging in our values for β, k and p, the first term on the right hand side above is $\exp\left(-\Theta(q^{-2c\ell/3})\right)$. Using $v \leqslant 50\sqrt{pk\log(1/\eta)} \leqslant q^{(c+o(1))\ell}$, the second term on the right hand side is at least $\exp\left(-\Omega(q^{(-c/3+o(1))\ell})\right)$. Finally, the last term is at least $\exp(q^{-2\ell})$. Overall, we get that for large enough ℓ

$$\frac{D'(x_1,\ldots,x_{2\ell})}{D(x_1,\ldots,x_{2\ell})} \geqslant \exp\left(-O(q^{-(c/3-o(1))\ell})\right) \geqslant \frac{1}{1.1}.$$

For the other direction, we show

$$\frac{D'(x_1,\ldots,x_{2\ell})}{D(x_1,\ldots,x_{2\ell})} \leqslant \frac{1}{0.9},$$

in nearly the same fashion. First note that

$$\left(\frac{1-\beta+\beta q^{2\ell}}{1-\beta}\right)^{s'} \leqslant \left(\frac{1-\beta+\beta q^{2\ell}}{1-\beta}\right)^{\ell^{100}} = \left(1+\frac{\beta}{1-\beta}q^{2\ell}\right)^{\ell^{100}} \leqslant 1+o(1). \tag{26}$$

By Equation (24), we have

$$\begin{aligned}
\frac{D'(x_1,\ldots,x_{2\ell})}{D(x_1,\ldots,x_{2\ell})} &= (1-\beta)^{k-s-s'}\left(1-\beta+\frac{\beta}{3}q^{2\ell}\right)^{s}(1-\beta+\beta q^{2\ell})^{s'} \\
&\leqslant (1+o(1))\cdot(1-\beta)^{k-s}\left(1-\beta+\frac{\beta}{3}q^{2(2\ell)}\right)^{s} \\
&= (1+o(1))\cdot(1-\beta)^{k-s}\left(1+\beta\left(\frac{q^{2\ell}}{3}-1\right)\right)^{s} \\
&\leqslant (1+o(1))\cdot\exp\left(-\beta(k-s)+\left(\frac{q^{2\ell}}{3}-1\right)\beta s\right) \\
&= (1+o(1))\cdot\exp\left(-\beta k+\frac{q^{2\ell}}{3}\beta s\right).
\end{aligned}$$

where in the first transition we use Equation (26), and in the fourth transition we use the fact that $1+z \leqslant e^z$. Writing $s = pk+v$ and using Equation (25) again and using $v \leqslant O\left(q^{(c/2+o(1))\ell}\right)$ we have,

$$\frac{D'(x_1,\ldots,x_{2\ell})}{D(x_1,\ldots,x_{2\ell})} \leqslant (1+o(1))\exp\left(-\frac{2}{3}\beta k q^{-2\ell}+\frac{q^{2\ell}}{3}\beta v\right) \leqslant (1+o(1))\exp\left(O(q^{-c\ell/3-o(1)})\right) \leqslant \frac{1}{0.9}.$$

D List Decoding Bound

In this section we prove Lemma 5.20, restated below.

Lemma 5.20. *Let T be a table on $\mathsf{Grass}_q(\mathbb{F}_q^V, 2\ell)$, let Q be an r_1-dimensional subspace, and let $W \supseteq Q$ be a subspace of codimension r_2. Suppose that 2ℓ is sufficiently large and $\dim(W) \geqslant 20\ell$. Let $f_1, \ldots, f_m$ be a list of distinct linear functions such that $f_i|_L \equiv T[L]$ for at least β-fraction of the 2ℓ-dimensional subspaces L such that $Q \subseteq L \subseteq W$, for $\beta \geqslant 2q^{-2\ell+r_1} + c$, and $c > 0$. Then,*

$$m \leqslant \frac{4}{c^2} \leqslant \frac{4}{\beta^2}.$$

Lemma 5.20 follows directly from a generic list decoding bound of [GRS00, Theorem 15], which we state below for convenience.

Theorem D.1. *[[GRS00, Theorem 15]] Let $\mathcal{C} \subseteq [Q]^N$ be a code with alphabet size Q, blocklength N, and relative distance $1-\gamma$. Let $\delta > 0$ and $R \in \Sigma^N$, where $Q = |\Sigma|$. Suppose that $C_1, \ldots, C_m \in \Sigma^N$ are distinct codewords from $\mathcal{C}$ that satisfy $\delta(R, C_i) \leqslant 1 - \delta$ for all $1 \leqslant i \leqslant m$. If,*

$$\delta > \frac{1}{Q} + \sqrt{\left(\gamma - \frac{1}{Q}\right)\left(1 - \frac{1}{q}\right)},$$

then

$$m \leqslant \frac{1}{(\delta - 1/Q)^2 - (1-1/Q)(\gamma - 1/Q)}.$$

Proving Lemma 5.20 is simply a matter of translating to the notation of Theorem D.1.

Proof of Lemma 5.20. Let $\Sigma = \mathbb{F}_q^{2\ell - r_1}$, let $Q = \mathrm{span}(z_1, \ldots, z_{r_1})$, and define a code $\mathcal{C}$ with alphabet Σ by:

$$\mathcal{C} = \{(v \cdot x_1, \ldots, v \cdot x_{2\ell - r_1})_{(x_1, \ldots, x_{2\ell - r_1}) \in W} \mid v \in W\}.$$

Note that for distinct $v, w \in W$ we have that $v \cdot x = w \cdot x$ for at most $1/q$-fraction of $x \in W$. Thus, the relative distance of C is $1 - q^{-2\ell + r_1}$. We would like the table T corresponds to a word, say, R, and $f_1, \ldots, f_m$ correspond to m codewords in $\mathcal{C}$, say $C_1, \ldots, C_m$. A slight issue is that T is only defined over 2ℓ-dimensional subspaces of $L \in \mathsf{Zoom}[Q, W]$, while R has an entry for every $2\ell - r_1$-tuple of points in V. To resolve this, note that nearly every $2\ell - r_1$-tuple of points combined with $z_1, \ldots, z_{r_1}$ span an $L \in \mathsf{Zoom}[Q, W]$. Thus, define R as follows. If $(z_1, \ldots, z_{r_1}, x_1, \ldots, x_{2\ell - r_1})$ are linearly independent, then let L be the span of $(z_1, \ldots, z_{r_1}, x_1, \ldots, x_{2\ell - r_1})$ and define

$$R_{(x_1, \ldots, x_{2\ell - r_1})} = (T[L](x_1), \ldots, T[L](x_{2\ell - r_1})).$$

Otherwise, define $R_{(x_1, \ldots, x_{2\ell - r_1})}$ arbitrarily. Note that the fraction of tuples $(x_1, \ldots, x_{2\ell - r_1})$ such that $(z_1, \ldots, z_{r_1}, x_1, \ldots, x_{2\ell - r_1})$ are not linearly independent is at most,

$$\sum_{i = r_1 + 1}^{2\ell} \frac{q^{i-1}}{q^n} \leqslant q^{2\ell - n},$$

so nearly all of the entries in R correspond to table entries in T. For the functions $f_1, \ldots, f_m$ we define C_i corresponding to f_i by

$$C_{i_{(x_1, \ldots, x_{2\ell - r_1})}} = (f_i(x_1), \ldots, f_i(x_{2\ell - r_1})).$$

As each f_i agrees with T on at least β-fraction of the entries, we have that R and C_i agree on at least β-fraction of the entries $(x_1, \ldots, x_{2\ell-r_1})$ such that $(z_1, \ldots, z_{r_1}, x_1, \ldots, x_{2\ell-r_1})$ are linearly independent, so

$$\delta(R, C_i) \leqslant 1 - \beta \cdot (1 - q^{2\ell-n}) \leqslant 1 - \frac{\beta}{2}$$

for each $1 \leqslant i \leqslant m$. Finally, note that the alphabet size of $\mathcal{C}$ is $\left|\mathbb{F}_q^{2\ell-r_1}\right| = q^{2\ell-r_1}$. To bound m, we can apply Theorem D.1 with $\delta = \frac{\beta}{2} \geqslant q^{-2\ell} + \frac{c}{2}$, $Q = q^{2\ell-r_1}$, and $\gamma = q^{-2\ell+r_1}$. We first note that the condition of Theorem D.1 is indeed satisfied,

$$\delta \geqslant q^{-2\ell+r_1} + \frac{c}{2} > q^{-2\ell+r_1} + 0.$$

Thus Theorem D.1 implies that $m \leqslant \frac{4}{c^2}$. $\square$

E Missing Proofs from Section 8

This section contains the missing proofs from Section 8, and we begin by recalling some notation. We denote by $\mu(\mathcal{A})$ the measure of a collcetion of subspaces $\mathcal{A} \subseteq \mathsf{Grass}_q(n, i)$, where n and i will always be clear from context. Furthermore, we use $\mu_X(\mathcal{A})$ to denote the measure of $\mathcal{A}$ restricted to the subspaces containing X for some subspace low-dimensional subspace X, i.e.

$$\mu_X(\mathcal{A}) = \frac{|\{L \in \mathcal{A} \mid X \subseteq L\}|}{|\{L \in \mathsf{Grass}_q(n, i) \mid X \subseteq L\}|}.$$

Likewise, when W is a constant co-dimension subspace, we define

$$\mu_W(\mathcal{A}) = \frac{|\{L \in \mathcal{A} \mid L \subseteq W\}|}{|\{L \in \mathsf{Grass}_q(n, i) \mid L \subseteq W\}|}.$$

It will always be clear from the size of X or W in context which of the above definitions we are referring to. We also use

$$\mu_{[X,W]}(\mathcal{A}) = \frac{|\{L \in \mathcal{A} \mid X \subseteq L \subseteq W\}|}{|\{L \in \mathsf{Grass}_q(n, i) \mid X \subseteq L \subseteq W\}|},$$

to denote the measure of $\mathcal{A}$ restricted to the zoom-in of X and the zoom-out of W. Finally, throughout this section, for some subspace L, and a set of constant codimension subspaces $\mathcal{W}$, we will let $N_{\mathcal{W}}(L) = |\{W \in \mathcal{W} \mid L \subseteq W\}|$.

E.1 Proof of Lemma 8.1

Recall that $\mathcal{W}$ is a set of m_1 subspaces of codimension s inside of V that is t-generic with respect to V. For each $2\left(1 - \frac{\xi}{2}\right)\ell$-dimensional subspace X and linear assignment, σ, to X, define

$$p_{X,\sigma} = \Pr_{\substack{W_i \in \mathcal{W}_{X,\sigma} \\ X \subseteq L \subseteq W_i}}[L \in \mathcal{L}_{X,\sigma} \wedge f_i|_L \neq T[L]], \qquad q_{X,\sigma} = \Pr_{\substack{W_i \in \mathcal{W}_{X,\sigma} \\ X \subseteq L \subset W_i}}[L \in \mathcal{L}_{X,\sigma}],$$

where in both probabilities X and σ are fixed, and $W_i \in \mathcal{W}_{X,\sigma}$ is chosen uniformly and $L \in \mathsf{Zoom}[X, W_i]$ is chosen uniformly. The intention behind these values is that for a fixed (X, σ), the quantity $p_{X,\sigma}$ should reflect how much disagreement there is between the table T and the functions f_i for $W_i \in \mathcal{W}_{X,\sigma}$, on subspaces $L \in \mathcal{L}_{X,\sigma}$, while $q_{X,\sigma}$ should reflect the size of $\mathcal{L}_{X,\sigma}$. Note that if $L \in \mathcal{L}_{X,\sigma}$ and $W_i \in \mathcal{W}_{X,\sigma}$, then by definition we already have $f_i|_X \equiv T[L]|_X \equiv \sigma$. Therefore we would expect that in fact f_i and $T[L]$ also agree on L - which is only larger than X by $\xi\ell$ dimensions - and for most X, σ, the value $p_{X,\sigma}$ is small. On the other hand, for each W_i, there are at least a C-fraction of $L \in \mathsf{Grass}_q(W_i, 2\ell)$ for which $f_i|_L \equiv T[L]$, so we should also expect $q_{X,\sigma}$ to be $\Omega(C)$ for a non-trivial fraction of (X, σ). In the following claim, we formalize this intuition and show that there indeed exists an (X, σ) for which $p_{X,\sigma}$ is small, $q_{X,\sigma}$ is large, and additionally the set $\mathcal{W}_{X,\sigma}$ is large.

This idea of looking for such (X, σ) was first introduced in [IKW12] where they call these (X, σ)-excellent and was then used again in [BDN17, MZ23] to analyze lower dimensional subspace versus subspace tests, which is similar in spirit to what we are ultimately trying to show in Lemma 5.19.

Claim E.1. *There exists (X, σ) and $\tau \geqslant \frac{C}{2}$ such that:*

- $m_2 = |\mathcal{W}_{X,\sigma}| \geqslant \frac{m_1}{q^{10r\ell}}$.
- $q_{X,\sigma} \geqslant \tau$.
- $p_{X,\sigma} < \gamma \cdot \tau$.

Proof. Consider the following process which outputs W_i, L, X, σ such that W_i is uniform in $\mathcal{W}$, $L \in \mathsf{Grass}_q(W_i, 2\ell)$ is uniform, $X \in \mathsf{Grass}_q\left(L, 2\left(1 - \frac{\xi}{2}\right)\ell\right)$ is uniform, and σ is the assignment of f_i to X, i.e $\sigma :\equiv f_i|_X$.

1. Choose (X, σ) with probability proportional to $|\mathcal{W}_{X,\sigma}|$.
2. Choose $W_i \in \mathcal{W}_{X,\sigma}$ uniformly.
3. Choose a 2ℓ-dimensional subspace L uniformly conditioned on $X \subseteq L \subseteq W$.

Notice that the distribution of (W_i, L) above is equivalent to that of choosing $W_i \in \mathcal{W}$ uniformly and $L \subset W_i$ uniformly. Moreover, $f_i|_L \equiv T[L]$ only if $L \in \mathcal{L}_{X,\sigma}$, as f_i and $T[L]$ must agree on $X \subseteq L$ in order to agree on L. Therefore,

$$\mathop{\mathbb{E}}_{X,\sigma}[q_{X,\sigma}] \geqslant \Pr_{W_i \in \mathcal{W}, L \subseteq W_i}[f_i|_L \equiv T[L]] \geqslant C \geqslant \frac{1}{q^{2(1-\xi)\ell}}. \tag{27}$$

On the other hand,

$$\mathop{\mathbb{E}}_{X,\sigma}[p_{X,\sigma}] \leqslant \Pr_{X \subseteq L}[f_i|_X \equiv T[L]|_X \ \mid \ f_i|_L \not\equiv T[L]] \leqslant \frac{1}{q^{2(1-\xi/2)\ell}}. \tag{28}$$

Here the distribution of (X, σ) is proportional to the sizes $|\mathcal{W}_{X,\sigma}|$ and the second inequality is by the Schwartz-Zippel lemma. Indeed, by the Schwartz-Zippel lemma, $f_i|_L$ and $T[L]$ can agree on at most $1/q$-fraction of points z in L. Therefore, the middle term is bounded by the probability that $2(1-\xi/2)\ell$ uniformly random, linearly independent points are all chosen in this $1/q$-fraction.

From a dyadic-partitioning of Equation (27), it follows that there exists a $\tau \geqslant C/2$ such that

$$\Pr_{X,\sigma}[q_{X,\sigma} \in [\tau, 2\tau)] \geqslant \frac{C}{4\tau \log(1/C)}. \tag{29}$$

By Markov's inequality on Equation (28)

$$\Pr_{X,\sigma}[p_{X,\sigma} \geqslant \gamma\tau] \leqslant \frac{1}{\gamma\tau q^{2(1-\xi/2)\ell}} \leqslant \frac{C}{8\tau\log(1/C)}, \tag{30}$$

It follows that for at least $\frac{C}{8t\log(1/C)}$-fraction of (X,σ)'s (under the measure induced by step 1 of the sampling procedure above), we have both $q_{X,\sigma} \geqslant \tau$ and $p_{X,\sigma} \leqslant \gamma\tau$.

Next we wish to argue that for most of these (X,σ)'s, $|\mathcal{W}_{X,\sigma}|$ is large. First note that the total number of (X,σ)'s is at most $\begin{bmatrix}\dim(V') \\ 2(1-\xi/2)\ell\end{bmatrix}_q q^{2\ell}$. For a fixed (X,σ), the probability that it is chosen is precisely,

$$\frac{|\mathcal{W}_{X,\sigma}|}{m} \cdot \frac{1}{\begin{bmatrix}\dim(V)-r \\ 2(1-\xi/2)\ell\end{bmatrix}_q}.$$

Thus, by a union bound,

$$\Pr_{X,\sigma}\left[|\mathcal{W}_{X,\sigma}| \leqslant \frac{m}{q^{10r\ell}}\right] \leqslant \frac{1}{q^{10r\ell}} \cdot \frac{\begin{bmatrix}n \\ 2(1-\xi/2)\ell\end{bmatrix}_q q^{2\ell}}{\begin{bmatrix}n-r \\ 2(1-\xi/2)\ell\end{bmatrix}_q} \leqslant \frac{1}{q^{10r\ell}} \cdot q^{4r\ell} \cdot q^{2\ell} \leqslant \frac{1}{q^{5r\ell}}. \tag{31}$$

Putting Equations (29), (30), and (31), together, it follows that with probability at least

$$\frac{C}{4\tau\log(1/C)} - \frac{1}{\gamma\tau q^{2(1-\xi/2)\ell}} - \frac{1}{q^{5r\ell}} > 0,$$

over (X,σ), we have, $q_{X,\sigma} \geqslant \tau$, $p_{X,\sigma} \leqslant \gamma\tau$, and $|\mathcal{W}_{X,\sigma}| \geqslant \frac{m_1}{q^{10r\ell}}$, which establishes the claim. □

Taking the (X,σ) given by Claim E.1, it almost looks like Lemma 8.1 is satisfied. However, notice that while the probability of interest for the third item there looks similar to $p_{X,\sigma}$, it has a different distribution over L and W_i. Indeed, there, the distribution first chooses $L \in \mathcal{L}_{X,\sigma}$, then $W_i \in \mathcal{W}_{X,\sigma}$, whereas for $p_{X,\sigma}$, we are first choosing $W_i \in \mathcal{W}_{X,\sigma}$, and not conditioning $L \subseteq W_i$ being in the set $\mathcal{L}_{X,\sigma}$. Intuitively, we expect something like the following to hold,

$$\Pr_{L\in\mathcal{L}_{X,\sigma}, W_i\in\mathcal{W}_{X,\sigma}}[f_i|_L \neq T[L] \mid W_i \supseteq L] = \Pr_{W_i\in\mathcal{W}_{X,\sigma}, L\subseteq W_i}[f_i|_L \neq T[L] \mid L \in \mathcal{L}_{X,\sigma}] = \frac{p_{X,\sigma}}{q_{X,\sigma}} \leqslant \frac{\gamma\cdot\tau}{\tau},$$

and be done. These equalities are not actually true however, so the bulk of the transition from Claim E.1 to Lemma 8.1 is in formalizing this chain of equalities and converting from the distribution of $p_{X,\sigma}$ to that required by the third item of Lemma 8.1 without losing too much.

Proof of Lemma 8.1. Fix an (X,σ) such that Claim E.1 holds and define $\mathcal{W}_{X,\sigma}$ and $\mathcal{L}_{X,\sigma}$ accordingly. Let $m_2 = |\mathcal{W}_{X,\sigma}| \geqslant \frac{m_1}{q^{10r\ell}}$. In order to lower bound the measure $\mu_X(\mathcal{L}_{X,\sigma})$, we will apply Lemma 5.14 on the collection of subspaces $\mathcal{W}_{X,\sigma}$ with parameters $j = 2\ell$, $a = 2\left(1-\frac{\xi}{2}\right)\ell$. Then the measure ν over $\mathsf{Zoom}[X,V]$ in Lemma 5.14 is precisely that obtained by choosing $W_i \in W_{X,\sigma}$ uniformly, and then $L \in \mathsf{Zoom}[X,W_i]$ uniformly. Thus ν is precisely the distribution used to define $q_{X,\sigma}$, so by Lemma 5.14

$$\mu_X(\mathcal{L}) \geqslant \nu(\mathcal{L}) - \frac{3q^{\frac{s}{2}\xi\ell}}{\sqrt{m_2}} = q_{X,\sigma} - \frac{3q^{\frac{s}{2}\xi\ell}}{\sqrt{m_2}} \geqslant \frac{\tau}{3} \geqslant \frac{C}{6},$$

and the first two conditions of Lemma 8.1 are satisfied.

To show the third condition, it will be helpful to have in mind the bipartite graph with parts $\mathcal{W}_{X,\sigma}$ and $\mathcal{L}_{X,\sigma}$ and edges (W_i, L) if $L \subseteq W_i$. For each $L \in \mathcal{L}_{X,\sigma}$ and $W_i \in \mathcal{W}_{X,\sigma}$ define the following degree-like quantities:

- $d_L = |\{W_i \in \mathcal{W}_{X,\sigma} \mid W_i \supseteq L\}|$,
- $e_L = |\{W_i \in \mathcal{W}_{X,\sigma} \mid W_i \supseteq L, f_i|_L \neq T_1[L]\}|$,
- $d_i = |\{L \in \mathcal{L}_{X,\sigma} \mid L \subseteq W_i\}|$,
- $e_i = |\{L \in \mathcal{L}_{X,\sigma} \mid L \subseteq W_i, f_i|_L \neq T_1[L]\}|$.

Also let $D = |\{L \mid X \subseteq L \subseteq W_i, \dim(L) = 2\ell\}|$, where the $W_i \in \mathcal{W}_{X,\sigma}$ is arbitrary (the value is the same regardless which we pick). Then $\mathbb{E}_{L \in \mathcal{L}_X}[d_L] = \frac{m_2 D}{|\mathcal{L}_X|}$ and the probability that we are interested in bounding can be expressed as:

$$\Pr_{L \in \mathcal{L}_{X,\sigma}, W_i \in \mathcal{W}_{X,\sigma}}[f_i|_L \neq T_1[L] \mid L \subseteq W_i] = \mathop{\mathbb{E}}_{L \in \mathcal{L}_{X,\sigma}}\left[\frac{e_L}{d_L}\right].$$

Since $q_{X,\sigma} \geqslant t$ and $p_{X,\sigma} \leqslant \gamma t$, we have

$$q_{X,\sigma} \cdot m_2 \cdot D = \sum_{L \in \mathcal{L}_{X,\sigma}} d_L = \sum_{W_i \in \mathcal{W}_{X,\sigma}} d_i \geqslant m_2 \cdot D\tau, \tag{32}$$

and

$$p_{X,\sigma} \cdot m_2 \cdot D = \sum_{L \in \mathcal{L}_{X,\sigma}} e_L = \sum_{W_i \in \mathcal{W}_{X,\sigma}} e_i \leqslant m_2 \cdot \gamma\tau D. \tag{33}$$

By Lemma 5.16 and the very loose bound $\frac{D}{|\mathcal{L}_X|} \approx \frac{1}{q^{\xi\ell/2}} \geqslant \frac{1}{q^\ell}$, we have

$$\Pr_{L \in \mathcal{L}_X}\left[d_L \leqslant 0.9 \mathop{\mathbb{E}}_{L}[d_L]\right] \leqslant \frac{q^\ell}{m_2}. \tag{34}$$

Using $\mathbb{E}_{L \in \mathcal{L}_X}[d_L] = \frac{m_2 D}{|\mathcal{L}_X|}$, we have

$$\begin{aligned}
\mathop{\mathbb{E}}_{L \in \mathcal{L}_{X,\sigma}}\left[\frac{e_L}{d_L}\right] &\leqslant \Pr_{L \in \mathcal{L}_{X,\sigma}}\left[d_L \leqslant 0.9\frac{m_2 D}{|\mathcal{L}_X|}\right] + \mathop{\mathbb{E}}_{L \in \mathcal{L}_{X,\sigma}}\left[\frac{e_L}{0.9 m_2 D/|\mathcal{L}_X|}\right] \\
&\leqslant \frac{q^\ell/m_2}{\Pr_{L \in \mathcal{L}_X}[L \in \mathcal{L}_{X,\sigma}]} + \mathop{\mathbb{E}}_{L \in \mathcal{L}_{X,\sigma}}\left[\frac{e_L}{0.9 m_2 D/|\mathcal{L}_X|}\right] \\
&\leqslant \frac{q^\ell/m_2}{\tau/3} + \mathop{\mathbb{E}}_{L \in \mathcal{L}_{X,\sigma}}\left[\frac{e_L|\mathcal{L}_X|}{0.9 m_2 D}\right] \\
&\leqslant \frac{3q^\ell}{m_2\tau} + \frac{m_2 \cdot \gamma\tau D}{0.9 m_2 D} \cdot \frac{|\mathcal{L}_X|}{|\mathcal{L}_{X,\sigma}|} \\
&\leqslant \frac{3q^\ell}{m_2\tau} + \frac{\gamma\tau}{0.9} \cdot \frac{3}{\tau} \\
&\leqslant 5\gamma,
\end{aligned}$$

where in the second transition we use Equation (34) and in the fourth transition we use Equation (33). □

E.2 Proof of Lemma 8.3

Take the $X, \sigma, \mathcal{L}'_{X,\sigma}$, and $\mathcal{W}_{X,\sigma}$ guaranteed by Corollary 8.2, and recall V' is the ambient space and $\delta_2 = \frac{\xi}{100}$. As this section is more involved, we restate Lemma 8.3 as well as its setting. Recall that for a zoom-in A and zoom-out B such that $X \subseteq A \subseteq B \subseteq V'$, we write $V' = A \oplus V_0$ and $B = A \oplus V^\star$, where $V^\star \subseteq V_0$. Now let $\mathcal{W}^\star_{[A,B]} = \{W_i^\star \mid \exists W_i \in \mathcal{W}_{X,\sigma} \text{ s.t } A \oplus W_i^\star = W_i \cap B\}$. It is clear that each $W_i^\star \in \mathcal{W}^\star_{[A,B]}$ is contained inside of some $W_i \in \mathcal{W}_{X,\sigma}$, so for each $W_i^\star$, we may define $f_i^\star \equiv f_i|_{W_i^\star}$.

Lemma 8.3. *There is a zoom-in A and a zoom-out B such that the following holds with the notation above. There exists a collection of subspaces $\mathcal{W}^\star = \{W_1^\star, \dots, W_{m_3}^\star\} \subseteq \mathcal{W}^\star_{[A,B]}$ of codimension s with respect to $V^\star$ such that:*

1. *For some $\ell' \geqslant \frac{\xi}{3}\ell$ there exists $\mathcal{L}^\star \subseteq \mathsf{Grass}_q(V^\star, \ell')$ such that $\mu(\mathcal{L}^\star) = \eta \geqslant \frac{C}{12}$.*
2. *The set $\mathcal{L}^\star$ is $(1, q^{\delta_2 \ell}\eta)$-pseudo-random.*
3. *Each $W_i^\star$ has codimension $s \leqslant r$ inside of $V^\star$ and $\mathcal{W}^\star$ is 4-generic, with respect to $V^\star$.*
4. *$m_3 \geqslant \frac{q^{-10s/\delta_2}}{2} \cdot m_2$.*
5. *For every $L \in \mathcal{L}^\star$, choosing $W_i^\star \in \mathcal{W}^\star$ uniformly such that $W_i^\star \supseteq L$, we have*
$$\Pr_{W_i^\star \supseteq L, W_i^\star \in \mathcal{W}^\star}[f_i|_L \neq T[L]] \leqslant 14\gamma.$$
6. *For every $L \in \mathcal{L}^\star$,*
$$0.8 \cdot m_3 \cdot q^{-s\cdot\ell'} \leqslant N_{\mathcal{W}^\star}(L) \leqslant 1.2 \cdot m_3 \cdot q^{-s\cdot\ell'},$$
where $N_{\mathcal{W}^\star}(L) = |\{W_i^\star \in \mathcal{W}^\star \mid W_i^\star \supseteq L\}|$

Here the table T is assigns linear functions to $L \in \mathsf{Grass}_q(V^\star, \ell)$, and is essentially the original table, i.e
$$T[L] \equiv T[A \oplus L]|_L.$$

As a step towards Lemma 8.3, we first show Lemma E.2, which finds the basic items required for Lemma 8.3, modulo a few minor alterations.

Lemma E.2. *We can find a zoom-in A and a zoom-out B such that $X \subseteq A \subseteq B \subseteq V'$, such that the following hold.*

- $\dim(A) + \operatorname{codim}(B) \leqslant \dim(X) + \frac{10}{\delta_2}$.
- *Letting $\mathcal{L}' = \{L \in \mathcal{L}'_{X,\sigma} \mid L \in \mathsf{Zoom}[A,B]\}$, we have $\eta = \mu_{[A,B]}(\mathcal{L}') \geqslant \frac{C}{12}$ in $\mathsf{Zoom}[A,B]$ and is $(1, q^{\delta_2\ell}\eta)$-global in $\mathsf{Zoom}[A,B]$.*[10]

Proof. Set $A_0 = X$, $\mathcal{L}_0 = \mathcal{L}_{X,\sigma}$, $B_0 = V$, and $\eta_0 = \mu_X(\mathcal{L}_{X,\sigma}) \geqslant \frac{C}{12}$. Now do the following.

1. Set $i = 0$, and initialize $A_0, \mathcal{L}_0, B_0, \eta_0, \mathcal{W}_{L,0}, \mathcal{W}_0$ as above.

[10] By $(1, q^{\delta_2\ell}\eta)$-pseudo-random in $\mathsf{Zoom}[A,B]$ we mean that $\mathcal{L}'$ does not increase its fractional size to $q^{\delta_2\ell}\eta$ when restricted to any zoom-in containing A or any zoom-out contained in B.

2. If $\mathcal{L}_i$ is $(1, q^{\delta_2 \ell}\eta_i)$-global inside of $\mathsf{Zoom}[A_i, B_i]$, then stop.

3. Otherwise, there exist $A \subseteq B$ such that, $A_i \subseteq A \subseteq B \subseteq B_i$, $\dim(A) + \operatorname{codim}(B) = \dim(A_i) + \operatorname{codim}(B_i) + 1$, and $\mu_{[A,B]}(\mathcal{L}_i) \geqslant q^{\delta_2 \ell}\eta_i$.

4. Set $A_{i+1} = A$, $B_{i+1} = B$, and $\mathcal{L}_{i+1} = \{L \in \mathcal{L}_{i+1} \mid A_{i+1} \subseteq L \subseteq B_{i+1}\}$.

5. Set $\eta_{i+1} = \mu_{[A_{i+1}, B_{i+1}]}(\mathcal{L}_{i+1})$.

6. Increment i by 1 and return to step 2.

Suppose this process terminates on iteration j. We claim that taking $\mathcal{L}' = \mathcal{L}_j$, $A = A_j$, $B = B_j$, $\mathcal{W}_L = \mathcal{W}_{L,j}$ for each $L \in \mathcal{L}_j$, and $\mathcal{W}' = \mathcal{W}_j$ satisfies the requirements of the lemma.

Next notice that by construction $\eta_{i+1} \geqslant q^{\delta_2 \ell}\eta_i$. Therefore, we perform at most $\frac{\log(12/C)}{\log(q^{\delta_2 \ell})} \leqslant \frac{10}{\delta_2}$ iterations before stopping, so $j \leqslant \frac{10}{\delta_2}$. By construction $\mathcal{L}_j$ is $(1, q^{\delta_2 \ell}\eta_j)$-global in $\mathsf{Zoom}[A_j, B_j]$ and has fractional size $\eta_j \geqslant \mu_{X,\sigma} \geqslant \frac{C}{12}$ in $\mathsf{Zoom}[A_j, B_j]$. Moreover, $\dim(A_j) + \operatorname{codim}(B_j) = \dim(X) + j \leqslant \dim(X) + \frac{10}{\delta_2}$, so the conditions of the lemma are satisfied. □

Take A, B and $\mathcal{L}'$ given by Lemma E.2. Before moving on the the straightforward derivation of Lemma 8.3, define

$$\mathcal{W}_{X,\sigma,L'} = \{W_i \in \mathcal{W}_{X,\sigma} \mid L' \subseteq W_i\} \quad \text{and} \quad \mathcal{W}_{\mathcal{L}'} = \bigcup_{L' \in \mathcal{L}'} \mathcal{W}_{X,\sigma,L'}.$$

Proof of Lemma 8.3. We now construct the $\mathcal{L}^\star, \mathcal{W}^\star$, and $V^\star$ that satisfy Lemma 8.3. Let $V^\star$ be a subspace such that $A \oplus V^\star = B$, set $\ell' = 2\ell - \dim(A)$, and let

$$\mathcal{L}^\star = \{L^\star \in \mathsf{Grass}_q(V^\star, \ell') \mid L^\star \oplus A \in \mathcal{L}'\}.$$

For each $L^\star \in \mathcal{L}^\star$, we will use L' to denote the corresponding subspace such that $A \oplus L = L' \in \mathcal{L}'$, and the fact that the correspondence,

$$L^\star \in \mathcal{L}^\star \longleftrightarrow L' = A \oplus L^\star \in \mathcal{L}',$$

is a bijection between $\mathcal{L}^\star$ and $\mathcal{L}'$. Recall that we abuse notation and let T to denote both the original table on 2ℓ-dimensional subspaces, as well as the new table on $\mathsf{Grass}_q(V^\star, \ell')$, given by $T[L] = T[L']|_L$. It will always be clear, based on the argument in $T[\cdot]$, which we are referring to. We obtain $\mathcal{W}^\star$ in a similar way as $\mathcal{L}^\star$, however, some care will be needed to ensure that it is 4-generic. First, set $\tilde{W}_{X,\sigma,\mathcal{L}'} = \{W_i \cap B \mid W_i \in \mathcal{W}_{X,\sigma,\mathcal{L}'}\}$, then take $\tilde{\mathcal{W}}^\star$ to be the maximal subset of $\tilde{W}_{X,\sigma,\mathcal{L}'}$ that is 4-generic with codimension s with respect to B. Finally, set

$$\mathcal{W}^\star = \{W_i^\star \subseteq V^\star \mid A \oplus W_i^\star \in \tilde{\mathcal{W}}^\star\},$$

and set $m_3 = |\mathcal{W}^\star|$. Summarizing, we have the following chain of relations,

$$\mathcal{W}_{X,\sigma} \xrightarrow{W_i \supseteq L' \in \mathcal{L}'} \mathcal{W}_{X,\sigma,\mathcal{L}'} \xleftrightarrow{\cap B} \tilde{W}_{X,\sigma,\mathcal{L}'} \xrightarrow{\text{Make 4-generic}} \tilde{\mathcal{W}}^\star \xleftrightarrow{\text{Subtract subspace } A} \mathcal{W}^\star \tag{35}$$

$$W_i \xrightarrow{W_i \supseteq L' \in \mathcal{L}'} W_i \xrightarrow{\cap B} W_i \cap B \longrightarrow W_i \cap B \longrightarrow W_i^\star \text{ s.t } W_i^\star \oplus A = W_i \cap B \tag{36}$$

which will be helpful to refer back to. The double arrow transitions are bijections, while in the single arrow transitions subspaces are being removed. The second line shows what a generic member of each set looks like, where W_i are the original subspaces in $\mathcal{W}_{X,\sigma}$. We remark that we allow $\tilde{W}_{X,\sigma,\mathcal{L}'}$ to be a multiset. For each $L^\star \in \mathcal{L}^\star$, we define

$$\mathcal{W}^\star_{L^\star} = \{W_i^\star \in \mathcal{W}^\star \mid L^\star \subseteq W_i^\star\}.$$

It is clear from Equation (35) that $\mathcal{W}^\star \subseteq \mathcal{W}_{[A,B]}$. We now verify that the six properties of Lemma 8.3 hold.

Property 1. The subspaces of $\mathcal{L}^\star$ are of dimension ℓ' inside $V^\star$, and

$$\ell' \geqslant 2\ell - \dim(A) \geqslant 2\ell - \dim(X) - \frac{10}{\delta_2} \geqslant \frac{\xi}{3}\ell.$$

Also, $\mu(\mathcal{L}^\star) = \eta$ is the same as the measure of $\mathcal{L}'$ inside $\mathsf{Zoom}[A, B]$ due to the bijection between $\mathcal{L}^\star$ and $\mathcal{L}'$. Therefore which is at least $\eta \geqslant \frac{C}{12}$ by the second part of Lemma E.2.

Property 2. Since $\mathcal{L}'$ does not increase its measure to $q^{\delta_2 \ell}\eta$ on any zoom-in containing A or zoom-out inside B, it follows that $\mathcal{L}^\star$ is $(1, q^{\delta_2\ell}\eta)$-pseudo-random.

Property 3. By construction, $\tilde{W}^\star$ is 4-generic inside of B. Since $B = A \oplus V^\star$, and all $W_i \in \tilde{W}^\star$ contain A, it follows that $\mathcal{W}^\star$ is 4-generic inside of $V^\star$.

We verify property 4 using properties 5 and 6, so we save it for last.

Property 6. Fix an $L^\star \in \mathcal{L}^\star$ with corresponding $L' \in \mathcal{L}' \subseteq \mathcal{L}_{X,\sigma}$. We will first show that for all $L^\star \in \mathcal{L}'$, the value $N_{\mathcal{W}^\star}(L^\star)$ is nearly the same, and in particular is nearly equal to $N_{\mathcal{W}_{X,\sigma}}(L')$. This has the secondary consequence that $m_3 \geqslant N_{\mathcal{W}^\star}(L^\star)$ is large. By Lemma 5.16 applied to the 4-generic set of subspaces $\mathcal{W}^\star$ and we can conclude that Property 5 holds for most (and in particular at least one) subspaces of $\mathcal{L}^\star$. We can then conclude that the same holds for all $L^\star \in \mathcal{L}^\star$.

Towards showing that all $N_{\mathcal{W}^\star}(L^\star)$ are nearly the same, note that

$$N_{\mathcal{W}^\star}(L^\star) = N_{\tilde{\mathcal{W}}^\star}(L') = N_{\tilde{\mathcal{W}}_{\mathcal{L}'}}(L') - \left|\tilde{\mathcal{W}}_{\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star\right|,$$

and additionally

$$N_{\tilde{\mathcal{W}}_{\mathcal{L}'}}(L') = N_{\mathcal{W}_{X,\sigma,\mathcal{L}'}}(L') = N_{\mathcal{W}_{X,\sigma}}(L'), \tag{37}$$

so

$$N_{\mathcal{W}^\star}(L^\star) = N_{\mathcal{W}_{X,\sigma}}(L') - \left|\tilde{\mathcal{W}}_{\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star\right|.$$

Since we already have bounds on $N_{\mathcal{W}_{X,\sigma}}(L')$ from the fourth part of Corollary 8.2, it is sufficient to upper bound $\left|\tilde{\mathcal{W}}_{\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star\right|$. Since $\mathcal{W}_{X,\sigma,\mathcal{L}'}$ is $2^{2+\frac{10}{\delta_2}}$-generic, by Lemma 5.13, the set $\tilde{\mathcal{W}}_{\mathcal{L}'}$ can be made 4-generic by removing at most $\mathrm{codim}(B) \cdot (2^{2+\frac{10}{\delta_2}}) \leqslant \frac{10}{\delta_2} \cdot 2^{2+\frac{10}{\delta_2}}$ of the subspaces $W_i \cap B$ in $\tilde{W}_{X,\sigma,\mathcal{L}'}$, so

$$\left|N_{\mathcal{W}^\star}(L^\star) - N_{\mathcal{W}_{X,\sigma}}(L')\right| \leqslant \left|\tilde{\mathcal{W}}_{\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star\right| \leqslant \frac{10}{\delta_2} \cdot 2^{2+\frac{10}{\delta_2}}, \qquad \forall L^\star \in \mathcal{W}^\star. \tag{38}$$

Using the fourth part of Corollary 8.2, we get the secondary consequence that

$$m_3 \geqslant 0.95 \cdot m_2 \cdot q^{\xi\ell \cdot s} - \frac{10}{\delta_2} \cdot 2^{2+\frac{10}{\delta_2}}.$$

Using Lemma 5.16, we can show that for at least one $L^\star$,

$$0.9 \cdot m_3 \cdot q^{-s\cdot\ell'} \leqslant N_{W^\star}(L^\star) \leqslant 1.1 \cdot m_3 \cdot q^{-s\cdot\ell'}, \tag{39}$$

and combining this with Equation (38) shows that Property 6 holds for all $L^\star \in \mathcal{L}^\star$.

Property 5. We encourage the reader to refer back to the chain of relations in Equation (35) for this part. At a high level, we will start with a probability regarding $W_i^\star \in \mathcal{W}^\star$ at the right end of the chain, and gradually move leftwards and relate this to a probability regarding $W_i \in \mathcal{W}_{X,\sigma}$ - which we have a bound on from the fourth item of Corollary 8.2. To start, note that

$$\begin{aligned}\Pr_{W_i^\star \in \mathcal{W}^\star}[f_i^\star|_{L^\star} \neq T[L^\star] \mid L^\star \subseteq W_i^\star] &= \Pr_{W_i^\star \in \mathcal{W}^\star}[f_i|_{L^\star} \neq T[L^\star] \mid L^\star \subseteq W_i^\star] \\ &\leqslant \Pr_{W_i \cap B \in \tilde{W}^\star}[f_i|_{L'} \neq T[L'] \mid L' \subseteq W_i \cap B].\end{aligned}$$

The first transition is simply due to the fact that $f_i^\star = f_i|_{W_i^\star}$. For the second transition we use the fact that there is a one-to-one correspondence between $W_i^\star \in \mathcal{W}^\star$ and $W_i = A \oplus W_i^\star \in \tilde{\mathcal{W}}^\star$. For this pair $W_i^\star$, W_i, we have $f_i^\star = f_i|_{W_i^\star}$. Therefore the W_i from the second probability can be sampled by first choosing $W_i^\star$ according to the first distribution of the first probability, and then outputting $A \oplus W_i^\star$. The second transition then follows.

Next we have,

$$\begin{aligned}&\Pr_{W_i \cap B \in \tilde{W}^\star}[f_i|_{L'} \neq T[L'] \mid L' \subseteq W_i \cap B] \\ &= \Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[f_i|_{L'} \neq T[L'] \mid L' \subseteq W_i \cap B, W_i \cap B \in \tilde{\mathcal{W}}^\star] \\ &= \frac{\Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[f_i|_{L'} \neq T[L'], W_i \cap B \in \tilde{\mathcal{W}}^\star \mid L' \subseteq W_i \cap B]}{\Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[W_i \cap B \in \tilde{\mathcal{W}}^\star \mid L' \subseteq W_i \cap B]} \\ &\leqslant \Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[f_i|_{L'} \neq T[L'] \mid L' \subseteq W_i \cap B] + 2 \Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star \mid L' \subseteq W_i \cap B],\end{aligned}$$

where both transitions rely on the fact that $\tilde{W}^\star \subseteq \tilde{W}_{X,\sigma,\mathcal{L}'}$. In the last transition, we used the fact that $1/(1-\delta) \leqslant 1 + 2\delta$ if $\delta \leqslant 1/2$, thus if $\Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star \mid L' \subseteq W_i \cap B] \leqslant 1/2$ then the inequality holds. Else, $2\Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star \mid L' \subseteq W_i \cap B] \geqslant 1$ and the expression on the third line is at most 1 so the inequality on the last transition holds trivially. We will now analyze the last two terms separately. The second term can be bounded as follows,

$$\begin{aligned}\Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star \mid L' \subseteq W_i \cap B] &= \frac{\left|\{W_i \cap B \in \tilde{\mathcal{W}}_{\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star \mid L' \subseteq W_i \cap B\}\right|}{|\{W_i \cap B \in \tilde{\mathcal{W}}_{\mathcal{L}'} \mid L' \subseteq W_i \cap B\}|} \\ &\leqslant \frac{\left|\tilde{\mathcal{W}}_{\mathcal{L}'} \setminus \tilde{\mathcal{W}}^\star\right|}{N_{\mathcal{W}_{X,\sigma}}(L')} \\ &\leqslant \gamma.\end{aligned}$$

The first transition is evident, for the second transition note that the numerator does not decrease, while the denominator is the same (it follows from $N_{\tilde{W}_{X,\sigma,\mathcal{L}'}} = N_{\mathcal{W}_{X,\sigma}}(L')$ in Equation (37)), finally the third transition uses Equation (38) and the fact that $N_{\mathcal{W}_{X,\sigma}}(L')$ is large from the fourth item of Corollary 8.2.

For the first term, note that,

$$\Pr_{W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}}[f_i|_{L'} \neq T[L'] \mid L' \subseteq W_i \cap B] = \Pr_{W_i \in \mathcal{W}_{X,\sigma,\mathcal{L}'}}[f_i|_{L'} \neq T[L'] \mid L' \subseteq W_i],$$

where we use the fact that there is a one-to-one correspondence between $W_i \cap B \in \tilde{W}_{X,\sigma,\mathcal{L}'}$, and $W_i \in \mathcal{W}_{X,\sigma,\mathcal{L}'}$. Also, recalling the definition of $\mathcal{L}'$, we have $L' \subseteq B$, so the conditioning in both probabilities is the same, and therefore the two probabilities are equivalent. Next, it is clear that $\{W_i \in \mathcal{W}_{X,\sigma,\mathcal{L}'} \mid L' \subseteq W'\} = \mathcal{W}_{X,\sigma,L'}$, so using Corollary 8.2, we have,

$$\Pr_{W_i \in \mathcal{W}_{X,\sigma,\mathcal{L}'}}[f_i|_{L'} \neq T[L'] \mid L' \subseteq W_i] \leqslant \Pr_{W_i \in \mathcal{W}_{X,\sigma,L'}}[f_i|_{L'} \neq T[L']] \leqslant 12\gamma.$$

Putting everything together, we get that

$$\Pr_{W_i^\star \in \mathcal{W}^\star}[f_i^\star|_{L^\star} \neq T[L^\star] \mid L^\star \subseteq W_i^\star] \leqslant 12\gamma + 2\gamma = 14\gamma,$$

establishing property 6.

Property 4. Take an arbitrary $L^\star \in \mathcal{L}^\star$ with corresponding $L' \in \mathcal{L}'$ such that Equation (39) holds. We have

$$\begin{aligned}
m_3 &\geqslant \frac{q^{-s\cdot\ell'}}{1.1} N_{\mathcal{W}^\star}(L^\star) \\
&\geqslant \frac{q^{-s\cdot\ell'}}{1.2}\left(N_{\mathcal{W}_{X,\sigma}}(L') - \frac{10}{\delta_2}\cdot 2^{1+\frac{10}{\delta_2}}\right) \\
&\geqslant \frac{q^{-s\cdot\ell'}}{1.2}\left(0.95\cdot m_2 \cdot q^{-\xi\ell\cdot s} - \frac{10}{\delta_2}\cdot 2^{1+\frac{10}{\delta_2}}\right) \\
&\geqslant \frac{m_2\cdot q^{s\cdot(\ell'-\xi\ell)}}{2} \\
&\geqslant \frac{m_2\cdot q^{-s\cdot(\dim(A)-\dim(X))}}{2} \\
&\geqslant \frac{m_2\cdot q^{-10s/\delta_2}}{2}.
\end{aligned}$$

The first transition is by Equation (39), the second transition is by Equation (37), the third transition is by the fourth item of Corollary 8.2, the fifth transition uses the fact that $\ell' = 2\ell - \dim(A)$ and $\dim(X) = 2\left(1-\frac{\xi}{2}\right)\ell$, and the last transition uses the fact that $\dim(A) - \dim(X) \leqslant \frac{10}{\delta_2}$. □

E.3 Proof of Lemma 8.5

Take $\mathcal{L}^\star, \mathcal{W}^\star, V^\star$ given by Lemma 8.3, and recall $Z = \{z \in V^\star \mid |\mu_z(\mathcal{L}^\star) - \eta| \leqslant \frac{\eta}{10}\}$. We have $\eta = \mu(\mathcal{L}^\star)$ and $m_3 = |\mathcal{W}^\star|$. Let us recall Lemma 8.5 below.

Lemma 8.5. *We have*

$$\Pr_{\substack{W_i^\star, W_j^\star \in \mathcal{W}^\star \\ z \in W_i^\star \cap W_j^\star \cap Z}}[f_i^\star(z) \neq f_j^\star(z)] \leqslant 500\gamma,$$

and for every $W_i^\star, W_j^\star \in \mathcal{W}^\star$,

$$|W_i^\star \cap W_j^\star \cap Z| \geqslant 0.81 \cdot |W_i^\star \cap W_j^\star|.$$

E.3.1 A Necessary Fourier Analytic Fact

We first show a Fourier Analytic fact that will be needed for the proof of Lemma 8.5. Let $\mathcal{A} \subseteq \mathsf{Grass}_q(n, j)$, for some $n > q^j$, and let $\eta = \mu(\mathcal{A})$. We define $F \in L_2\left(\mathbb{F}_q^{n\times j}\right)$ as follows:

$$F(x_1, \ldots, x_j) = \begin{cases} 1, & \text{if } \operatorname{span}(x_1, \ldots, x_j) \in \mathcal{A}, \\ 0, & \text{otherwise.} \end{cases}$$

Lemma E.3. *Fix a subspace $W \subseteq \mathbb{F}_q^n$, then for any $S = (s_1, \ldots, s_j) \in \mathbb{F}_q^{n\times j}$ we have,*

$$\mathop{\mathbb{E}}_{x_1,\ldots,x_j\in W}[\chi_S(x_1, \ldots, x_j)] = \begin{cases} 1, & \textit{if } S \subseteq W^\perp, \\ 0, & \textit{if } S \subsetneq W^\perp. \end{cases}$$

Proof. If $S \subseteq W^\perp$, then for any $x \in W$, we have $s_i \cdot x = 0$ for all $1 \leqslant i \leqslant j$, so the first case follows.

Now suppose $S \subsetneq W^\perp$, and without loss of generality say that $s_1 \notin W^\perp$. We can write,

$$\mathop{\mathbb{E}}_{x\subseteq W}[\chi_S(x)] = \mathop{\mathbb{E}}_{x_1\in W}\left[\omega^{\mathrm{Tr}(x_1\cdot s_1)}\right] \mathop{\mathbb{E}}_{x_2,\ldots,x_j\in W}\left[\omega^{\sum_{i=2}^{j}\mathrm{Tr}(x_i\cdot s_i)}\right].$$

We will show that $\mathbb{E}_{x_1\in W}\left[\omega^{\mathrm{Tr}(x_1\cdot s_1)}\right] = 0$. Notice that it is sufficient to show that $x_1 \cdot s_1$ takes each value in $\mathbb{F}_q$ with equal probability over uniformly random $x_1 \in W$. First, since $s_1 \notin W$, $\Pr_{x_1\in W}[x_1 \cdot s_1 = 0] = \frac{1}{q}$. Next note for any $\alpha \neq 0$,

$$\Pr_{x_1\in W}[x_1 \cdot s_1 = 1] = \Pr_{x_1\in W}[(\alpha x_1) \cdot s_1 = \alpha] = \Pr_{x_1\in W}[x_1 \cdot s_1 = \alpha].$$

Therefore, $x_1 \cdot s_1$ takes each of the $q-1$ nonzero values in $\mathbb{F}_q$ with probability $\frac{1}{q}$ over uniform $x_1 \in W$, and this concludes the proof. □

Lemma E.4. *If $W \subseteq \mathbb{F}_q^n$ has codimension s and satisfies,*

$$|\mu_W(\mathcal{A}) - \eta| \geqslant 0.01\eta,$$

then there is a nonzero $S \in \mathbb{F}_q^{n\times j}$ such that $S \subseteq W^\perp$ and

$$\left|\widehat{F}(S)\right| \geqslant \frac{\eta}{20q^{sj}}.$$

Proof. Note that, $\mu_W(\mathcal{A}) = \Pr_{x_1,\ldots,x_j}[\operatorname{span}(x_1, \ldots, x_j) \in \mathcal{A} \mid \dim(\operatorname{span}(x)) = j]$, so

$$\left|\mu_W(\mathcal{A}) - \mathop{\mathbb{E}}_{x\subseteq W}[F(x)]\right| \leqslant j \cdot q^{j-k}.$$

Using the Fourier decomposition of F, we can write,

$$\mathop{\mathbb{E}}_{x\subseteq W}[F(x)] = \widehat{F}(0) + \sum_{0\neq S\subseteq W^\perp} \widehat{F}(S) \mathop{\mathbb{E}}_{x\in W}[\chi_S(x)] + \sum_{\emptyset\neq S\subsetneq W^\perp} \widehat{F}(S) \mathop{\mathbb{E}}_{x\in W}[\chi_S(x)].$$

Using the previous inequality, and the fact that $\widehat{F}(0) = \eta$, and Lemma E.3, we have

$$\left| \mu_W(\mathcal{A}) - \eta - \sum_{0 \neq S \subseteq W^\perp} \widehat{F}(S) \right| \leqslant j \cdot q^{j-n}.$$

By the triangle inequality we have,

$$|\mu_W(\mathcal{A}) - \eta| \leqslant \left| \sum_{0 \neq S \subseteq W^\perp} \widehat{F}(S) \right| + j \cdot q^{j-k},$$

and finally by the assumption in the lemma statement we have,

$$\left| \sum_{0 \neq S \subseteq W^\perp} \widehat{F}(S) \right| \geqslant 0.1 \cdot \eta - j \cdot q^{j-k}.$$

Since $j \cdot q^{j-k} \leqslant 0.01 \cdot \eta$ and there are at most q^{sj} tuples $S = (s_1, \ldots, s_j) \subseteq W^\perp$, the result follows. □

E.3.2 The Proof of Lemma 8.5

For an arbitrary fixed point $z \in V^\star$, let D denote the number of ℓ'-dimensional subspaces $L \subseteq V^\star$ containing z. We note that D does not depend on which point z is fixed. Also let,

$$\begin{aligned}
&\mathcal{L}_z^\star = \{L \in \mathcal{L}^\star \mid z \in L\}, \\
&\mathcal{W}_z^\star = \{W_i^\star \in \mathcal{W}^\star \mid z \in W_i^\star\}, \\
&m_z = |\mathcal{W}_z^\star|, \\
&N_{2,\mathcal{W}_z^\star}(L) = |\{(i,j) \mid W_i^\star \cap W_j^\star \supseteq L,\ W_i^\star, W_j^\star \in \mathcal{W}_z^\star\}|, \\
&N_{2,\mathcal{W}^\star}(L) = |\{(i,j) \mid W_i^\star \cap W_j^\star \supseteq L, W_i^\star, W_j^\star \in \mathcal{W}^\star\}|.
\end{aligned} \tag{40}$$

Also for an arbitrary $W_i^\star$ and $W_j^\star$, define

$$p_1 = \Pr_{L \in \mathsf{Grass}_q(V^\star, \ell')}[L \subseteq W_i^\star] \quad \text{and} \quad p_2 = \Pr_{L \in \mathsf{Grass}_q(V^\star, \ell')}[L \in W_i^\star \cap W_j^\star \mid z \in L]. \tag{41}$$

A straightforward computation shows that $p_2/p_1^2 \geqslant q^{2s}/2$, where recall $s = \mathrm{codim}(W_i^\star)$ in $V^\star$. We start by removing all $z \in Z$ that do not satisfy,

$$1.1 \cdot q^{-s} m_3 \geqslant m_z \geqslant 0.9 \cdot q^{-s} m_3. \tag{42}$$

By Lemma 5.16, the number of z removed is at most $\frac{2q^s}{m_3}$ and is negligible, so for the remainder of the section we assume that all $z \in Z$ satisfy the above inequalities.

We now define two distributions that we will later show are close to each other. The first is $\mathcal{D}_1$, generated by choosing $z \in Z$ uniformly and $W_i^\star, W_j^\star \in \mathcal{W}^\star$ uniformly conditioned on $z \in W_i^\star \cap W_j^\star$. The second is $\mathcal{D}_1'$, generated by choosing $z \in Z$ uniformly, $L \in \mathcal{L}_z^\star$ uniformly, and then $W_i^\star, W_j^\star \in \mathcal{W}^\star$ uniformly conditioned on $L \subseteq W_i^\star \cap W_j^\star$. We have

$$\begin{aligned}
\mathcal{D}_1(z, W_i^\star, W_j^\star) &= \frac{1}{|Z|} \cdot \frac{1}{|\{i, j \mid W_i^\star, W_j^\star \in \mathcal{W}^\star, z \in W_i^\star \cap W_j^\star\}|}, \\
\mathcal{D}_1'(z, W_i^\star, W_j^\star) &= \frac{1}{|Z|} \cdot \frac{|\{L \in \mathcal{L}_z^\star \mid L \subseteq W_i^\star \cap W_j^\star\}|}{|\mathcal{L}_z^\star|} \cdot \mathop{\mathbb{E}}_{L \in \mathcal{L}_z^\star, L \subseteq W_i^\star \cap W_j^\star}\left[\frac{1}{N_{2,\mathcal{W}^\star}(L)}\right].
\end{aligned} \tag{43}$$

By construction of Z in Lemma 8.4, we have $|\mathcal{L}_z^\star| \leqslant 1.1\eta \cdot D$ for all $z \in Z$. Also, since $N_{2,\mathcal{W}^\star}(L) = N_{\mathcal{W}^\star}(L)^2$, the fifth property of Lemma 8.3 yields $1.21 \cdot m_3^2 p_1^2 \geqslant N_{2,\mathcal{W}^\star}(L) \geqslant 0.81 \cdot m_3^2 p_1^2$, for all $L \in \mathcal{L}^\star$. Now, noting that $p_1 \geqslant q^{-2s}$, we have:

$$\frac{1}{|Z|} \cdot \frac{1}{1.21 m_3^2 q^{-2s}} \leqslant \mathcal{D}_1(z, W_i^\star, W_j^\star) = \frac{1}{|Z|} \cdot \frac{1}{|\{i, j \mid W_i^\star, W_j^\star \in \mathcal{W}^\star, z \in W_i^\star \cap W_j^\star\}|}$$
$$\leqslant \frac{1}{|Z|} \cdot \frac{1}{0.81 \cdot m_3^2 q^{-2s}}, \tag{44}$$

and

$$\mathcal{D}_1'(z, W_i^\star, W_j^\star) \geqslant \frac{1}{|Z|} \cdot \frac{\mu_{[z, W_i^\star \cap W_j^\star]}(\mathcal{L}^\star) \cdot D_u \cdot p_2}{1.1 \cdot \eta \cdot D_u} \cdot \frac{1}{1.21 \cdot m_3^2 p_1^2} \geqslant \frac{1}{|Z|} \cdot \frac{\mu_{[z, W_i^\star \cap W_j^\star]}(\mathcal{L}^\star)}{5 \cdot \eta \cdot m_3^2 q^{-2s}}. \tag{45}$$

By construction, $|\mu_z(\mathcal{L}_z^\star) - \eta| \leqslant \frac{\eta}{10}$ for every $z \in Z$. Call a triplet $(z, W_i^\star, W_j^\star)$ bad if

$$\mu_{[z, W_i^\star \cap W_j^\star]}(\mathcal{L}^\star) \leqslant \frac{4}{5}\eta.$$

If the triplet $(z, W_i^\star, W_j^\star)$ is not bad, then by the above inequalities

$$\mathcal{D}_1(z, W_i^\star, W_j^\star) \leqslant \frac{1}{|Z|} \cdot \frac{1}{0.81 \cdot m^2 q^{-2s}} \leqslant \frac{1}{|Z|} \cdot \frac{24\eta}{25 \cdot \eta \cdot m^2 q^{-2s}} \leqslant 6\mathcal{D}_1'(z, W_i^\star, W_j^\star). \tag{46}$$

We start by showing that there are very few bad triplets.

Claim E.5. *For each z, the number of i, j such that*

$$\mu_{[z, W_i^\star \cap W_j^\star]}(\mathcal{L}^\star) \leqslant \frac{4}{5}\eta,$$

is at most $\frac{10^6 q^{4s\ell}}{\eta^2} m_3$. Additionally, for every $W_i^\star, W_j^\star \in \mathcal{W}^\star$, we have

$$\mu_{W_i^\star \cap W_j^\star}(Z) \geqslant 0.9\mu(Z).$$

Proof. Fix a point z, let F_z be the restriction of F to the zoom-in of z, where $F(x_1, \dots, x_{\ell'}) = 1$ if $\mathrm{span}(x_1, \dots, x_{\ell'}) \in \mathcal{L}^\star$ and 0 otherwise. Let $\eta' = \mu_z(\mathcal{L}^\star)$. For any i, j satisfying the inequality of the lemma,

$$\left| \mu_{[z, W_i^\star \cap W_j^\star]}(\mathcal{L}^\star) - \eta' \right| \geqslant \frac{\eta'}{20}.$$

We can then apply Lemma E.4 to the zoom-in on z. By Lemma E.4, if $(z, W_i^\star, W_j^\star)$ is bad then there must be $S = (s_1, \dots, s_{\ell'-1})$ such that $\mathrm{span}(s_1, \dots, s_{\ell'-1}) \subseteq \left(W_i^\star \cap W_j^\star\right)^\perp$ and,

$$|\widehat{F_z}(S)| \geqslant \frac{\eta'}{400 q^{2s(\ell'-1)}}. \tag{47}$$

Since by Parseval's inequality the sum of $\left|\widehat{F_z}(S)\right|^2$ is at most $\|F_z\|_2^2 \leqslant 1$, there are at most $\frac{160000 q^{4s\ell'}}{\eta'^2}$ tuples S satisfying (47). Now consider a bipartite graph where the left side consists of these tuples $S = (s_1, \dots, s_{\ell'-1})$, the right side consists of $W_i^\star \cap W_j^\star$, and the edges are between pairs that satisfy

$$\mathrm{span}(s_1, \dots, s_{\ell'-1}) \subseteq \left(W_i^\star \cap W_j^\star\right)^\perp.$$

It follows that the number of edges in this graph is an upper bound on the number of bad triples containing z. Since $\mathcal{W}^\star$ is 4-generic, we have

$$(W_i^\star \cap W_j^\star)^\perp \cap (W_{i'}^\star \cap W_{j'}^\star)^\perp = \{0\}$$

for all i, j, i', j' distinct. Therefore, any two neighbours of a vertex on the left must either have their i or j be equal, and hence the maximum degree of a vertex on the left side is at most $2m_3$. As a result, the graph has at most $2m_3 \cdot \frac{160000q^{4s\ell}}{\eta^2} \leqslant \frac{10^6 q^{4s\ell}}{\eta'^2} m_3$ edges, where we also use that $\eta' \geqslant 0.9\eta$. This completes the proof of the first assertion of the claim.

For the second part of the lemma, note that $\mu(\overline{Z}) \leqslant \frac{q^{-\ell'}}{2}$ by Lemma 8.4. Therefore, for any $W_i^\star \cap W_j^\star$, we have,

$$\mu_{W_i^\star \cap W_j^\star}(\overline{Z}) \leqslant q^{2s} \cdot \frac{q^{-\ell'}}{2}.$$

It follows that,

$$\mu_{W_i^\star \cap W_j^\star}(Z) \geqslant 1 - q^{2s} \cdot \frac{q^{-\ell'}}{2} \geqslant 0.9\mu(Z).$$

□

Lemma E.6. *Let E be any event defined with respect to $(z, W_i^\star, W_j^\star)$. Then,*

$$\mathcal{D}_1(E) \leqslant 6\mathcal{D}_1'(E) + \gamma.$$

Proof. If the triple $(z, W_i^\star, W_j^\star)$ is not bad, then $\mathcal{D}_1(z, W_i^\star, W_j^\star) \leqslant 6\mathcal{D}_1'(z, W_i^\star, W_j^\star)$. Otherwise, we can use the generic bound $\mathcal{D}_1(z, W_i^\star, W_j^\star) \leqslant \frac{1}{0.81 \cdot m_3^2 q^{-2s}}$, which can be obtained from Equation (44). By the bound on the number of bad triples per z in Claim E.5, it follows that

$$\begin{aligned}
\mathcal{D}_1(E) &\leqslant 6\mathcal{D}_1'(E) + |Z| \cdot \frac{10^6 q^{4s\ell}}{\eta^2} m_3 \frac{1}{|Z| \cdot 0.81 \cdot m_3^2 q^{-2s}} \\
&= 6\mathcal{D}_1'(E) + \frac{10^7 q^{4s\ell+2s}}{\eta^2 \cdot m_3} \\
&\leqslant 6\mathcal{D}_1'(E) + \gamma.
\end{aligned}$$

Note that in the last transition we are using the fact that m_3 is large by the fourth property in Lemma 8.3. □

Now let $\mathcal{D}_2$ be the distribution obtained by choosing $W_i^\star, W_j^\star \in \mathcal{W}$ uniformly, and then choosing $z \in W_i^\star \cap W_j^\star \cap Z$ uniformly. We have

$$\mathcal{D}_2(z, W_i^\star, W_j^\star) = \frac{1}{m_3^2} \cdot \frac{1}{|W_i^\star \cap W_j^\star \cap Z|}.$$

Using essentially the same proof, we get the following lemma.

Lemma E.7. *Let E be any event defined with respect to $(z, W_i^\star, W_j^\star)$. Then,*

$$\mathcal{D}_2(E) \leqslant 2\mathcal{D}_1(E).$$

Proof. We apply Claim E.4 where $V = V^\star$ and $\mathcal{L} = Z$. By Claim E.5, we have $\mu_{W_i^\star \cap W_j^\star}(Z) > 0.9\mu(Z)$ for all i, j, or equivalently

$$|W_i^\star \cap W_j^\star \cap Z| \geqslant 0.9 \cdot |Z| q^{-2s}.$$

Thus, for all i, j and all z,

$$\mathcal{D}_2(z, W_i^\star, W_j^\star) \leqslant \frac{1}{0.9|Z| \cdot q^{-2s}} \cdot \frac{1}{m_3^2} \leqslant \frac{2}{1.21|Z|} \cdot \frac{1}{m_3^2 q^{-2s}} \leqslant 2D_1(z, W_i^\star, W_j^\star),$$

where we use Equation (44) for the third transition. It follows that

$$\mathcal{D}_2(E) \leqslant 2\mathcal{D}_1(E).$$

□

We are now ready to prove Lemma 8.5.

Proof of Lemma 8.5. By the sixth property in Lemma 8.3, for every $L \in \mathcal{L}^\star$, we have

$$\Pr_{W_i^\star \supseteq L, W_i^\star \in \mathcal{W}^\star}[f_i|_L \neq T_1[L]] \leqslant 14\gamma.$$

Let E denote the event over $(z, W_i^\star, W_j^\star)$ that $f_i(z) \neq f_j(z)$. It follows that,

$$\mathcal{D}_1(E) \leqslant \Pr_{\substack{W_i^\star, W_j^\star \supseteq L \\ W_i^\star, W_j^\star \in \mathcal{W}^\star}}[f_i|_L \neq T_1[L] \ \vee \ f_j|_L \neq T_1[L]] \leqslant 2 \cdot \Pr_{W_i^\star \supseteq L, W_i^\star \in \mathcal{W}^\star}[f_i|_L \neq T_1[L]] \leqslant 28\gamma.$$

Putting Lemmas E.6 and E.7 together,

$$\mathcal{D}_2(E) \leqslant 2(6\mathcal{D}_1(E) + \gamma) = 338\gamma,$$

proving the first part of Lemma 8.5.

For the second part, recall from the second part of Claim E.5 that for every pair i, j, we have

$$\left|W_i^\star \cap W_j^\star \cap Z\right| \geqslant 0.9 \cdot |Z| \, q^{-2s} \geqslant 0.81 \, |V^\star| \, q^{-2s} = 0.81 \cdot \left|W_i^\star \cap W_j^\star\right|.$$

□

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il