Hardness of Max-2Lin and Max-3Lin over integers, reals, and large cyclic groups

Ryan O'Donnell*
Department of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: odonnell@cs.cmu.edu

Yi Wu†
Theory Group
IBM Almaden Research
San Jose, CA 95120, USA
Email: wuyi@us.ibm.com

Yuan Zhou
Department of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213, USA
Email: yuanzhou@cs.cmu.edu

***Abstract*—In 1997, Håstad showed NP-hardness of $(1-\epsilon, 1/q+\delta)$-approximating Max-3Lin($\mathbb{Z}_q$); however it was not until 2007 that Guruswami and Raghavendra were able to show NP-hardness of $(1-\epsilon,\delta)$-approximating Max-3Lin($\mathbb{Z}$). In 2004, Khot–Kindler–Mossel–O'Donnell showed UG-hardness of $(1-\epsilon,\delta)$-approximating Max-2Lin($\mathbb{Z}_q$) for $q = q(\epsilon,\delta)$ a sufficiently large constant; however achieving the same hardness for Max-2Lin($\mathbb{Z}$) was given as an open problem in Raghavendra's 2009 thesis.**

In this work we show that fairly simple modifications to the proofs of the Max-3Lin($\mathbb{Z}_q$) and Max-2Lin($\mathbb{Z}_q$) results yield optimal hardness results over $\mathbb{Z}$. In fact, we show a kind of "bicriteria" hardness: even when there is a $(1-\epsilon)$-good solution over $\mathbb{Z}$, it is hard for an algorithm to find a δ-good solution over $\mathbb{Z}$, $\mathbb{R}$, or $\mathbb{Z}_m$ for any $m \geq q(\epsilon,\delta)$ of the algorithm's choosing.

I. Introduction

In this paper we consider one of the most fundamental algorithmic tasks: solving systems of linear equations. Given a ring R, the Max-kLin(R) problem is defined as follows: An input instance is a list of linear equations of the form $a_1 x_{i_1} + \cdots + a_k x_{i_k} = b$, where $a_1, \ldots, a_k, b \in R$ are constants and $x_{i_1}, \ldots, x_{i_k}$ are variables from the set $\{x_1, \ldots, x_n\}$. Each equation also comes with a nonnegative rational weight; it is assumed the weights sum up to 1. The algorithmic task is to assign values from R to the variables so as to maximize the total weight of satisfied equations. We say that an assignment is γ-*good* if the equations it satisfies have total weight at least γ. We say that an algorithm achieves (c,s)-*approximation* if, whenever the instance has a c-good solution, the algorithm is guaranteed to find an s-good solution.

*Part of this research was performed while the author was a member of the School of Mathematics, Institute for Advanced Study. Supported by NSF grants CCF-0747250, CCF-0915893, and DMS-0635607, BSF grant 2008477, and Sloan and Okawa fellowships. This material is based upon work supported by the National Science Foundation. Any opinions, finidings and conclusions or recommendations expressed in this material are those of the authors and do not necessarily reflect the views of the National Science Foundation.

†Part of this research was performed while the author was at Carnegie Mellon University.

A. Prior work on Max-3Lin($\mathbb{Z}$)

It is an old result of Arora–Babai–Stern–Sweedyk [ABSS93] that for all $0 < \delta < 1$ there exists $\epsilon > 0$ and $k \in \mathbb{Z}^+$ such that it is NP-hard to $(\epsilon, \delta\epsilon)$-approximate Max-kLin($\mathbb{Q}$). Håstad's seminal work from 1997 [Hås01] showed hardness even for very sparse, near-satisfiable instances: specifically, he showed that for all constant $\epsilon, \delta > 0$ and $q \in \mathbb{N}$, it is NP-hard to $(1-\epsilon, 1/q+\delta)$-approximate Max-3Lin($\mathbb{Z}_q$). This is optimal in the sense that it is algorithmically easy to $(1,1)$-approximate or $(c, 1/q)$-approximate Max-3Lin($\mathbb{Z}_q$). Håstad's hardness result even holds for the special case of Max-Γ-3Lin($\mathbb{Z}_q$), meaning that all equations are of the form $x_{i_1} - x_{i_2} + x_{i_3} = b$.

Håstad's proof does not strictly generalize the ABSS [ABSS93] result on Max-kLin($\mathbb{Q}$) because there is no obvious reduction from hardness over $\mathbb{Z}_q$ to hardness over $\mathbb{Q}$. Indeed, it was not until much later, 2006, that NP-hardness of $(1-\epsilon,\delta)$-approximating Max-kLin($\mathbb{Q}$) was shown [FGKP06], [GR06]. Finally, in 2007 Guruswami and Raghavendra [GR07] generalized all of [ABSS93], [FGKP06], [GR06] by showing NP-hardness of $(1-\epsilon,\delta)$-approximating Max-Γ-3Lin($\mathbb{Z}$). As we will see shortly, this easily implies the same hardness for Max-Γ-3Lin($\mathbb{Q}$) and Max-Γ-3Lin($\mathbb{R}$). Indeed, it shows a kind of "bicriteria" hardness: given a Max-Γ-3Lin($\mathbb{Z}$) instance with a $(1-\epsilon)$-good solution over $\mathbb{Z}$, it is NP-hard to find a δ-good solution even over $\mathbb{R}$. Guruswami and Raghavendra's proof followed that of Håstad's to some extent, but involved somewhat technically intricate derandomized Long Code testing, using Fourier analysis with respect to a certain exponential distribution on $\mathbb{Z}^+$.

We would also like to mention the very recent work of Khot and Moshkovitz [**?**]. Motivated by proving the Unique Games Conjecture, they showed a strong NP-hardness result for a homogeneous variant of Max-3Lin($\mathbb{R}$). Specifically, they considered the case where all equations are of the form $a_1 x_{i_1} + a_2 x_{i_2} + a_3 x_{i_3} = 0$ with $a_1, a_2, a_3 \in [\frac{1}{2}, 2]$. Very roughly

speaking, they showed there is a universal $\delta > 0$ such that for all $\epsilon > 0$ the following problem is NP-hard: given an instance where there is a "Gaussian-distributed" real assignment which is $(1-\epsilon)$-good, find a Gaussian-distributed assignment in which the weight of equations satisfied to within margin $\delta\sqrt{\epsilon}$ is at least $1-\delta$. This result is incomparable to the one in [GR07].

B. Prior work on Max-2Lin

Following Håstad's work there was five years of no progress on Max-2Lin(R) for any ring R. To circumvent this, in 2002 Khot [Kho02] introduced the Unique Games (UG) Conjecture, which would prove to be very influential (and notorious!). Khot showed a strong "UG-hardness" result for Max-2Lin$(\mathbb{Z}_2)$ (crediting the result essentially to Håstad), namely that for all $t > 1/2$ and sufficiently small $\epsilon > 0$ it is UG-hard to $(1-\epsilon, 1-\epsilon^t)$-approximate Max-2Lin$(\mathbb{Z}_2)$. This result is essentially optimal due to the Goemans–Williamson algorithm [GW95].

In 2004, Khot–Kindler–Mossel–O'Donnell [KKMO07] (using [MOO05]) extended this work by showing that for all $\epsilon, \delta > 0$, there exists $q \in \mathbb{N}$ such that $(1-\epsilon, \delta)$-approximating Max-Γ-2Lin$(\mathbb{Z}_q)$ is UG-hard, and hence in fact UG-*complete*. Here Γ-2Lin means that all equations are of the form $x_{i_1} - x_{i_2} = b$. KKMO gave a quantitative dependence as well: given ϵ and q one can choose any $\delta > q\Lambda_{1-\epsilon}(1/q) \approx (1/q)^{\epsilon/(2-\epsilon)}$, where $\Lambda_{1-\epsilon}(1/q)$ is a certain correlated Gaussian quadrant probability.

The following natural question was left open by KKMO [KKMO07]:

Question.: Is it true that for all $\epsilon, \delta > 0$ it is UG-*hard to $(1-\epsilon, \delta)$-approximate* Max-Γ-2Lin$(\mathbb{Z})$*?*

The key technical tool used in the KKMO hardness result for Max-2Lin$(\mathbb{Z}_q)$, namely the Majority Is Stablest Theorem [MOO05], has a bad dependence on the parameter q. Thus pushing q to be "superconstantly" large seemed to pose a fundamental problem. The question above is one of the open problems posed at the end of Raghavendra's monumental thesis [Rag09].

C. Our results

In this paper we show that it is relatively easy to modify the proofs of the hardness results known for Max-Γ-2Lin$(\mathbb{Z}_q)$ and Max-Γ-3Lin$(\mathbb{Z}_q)$ to obtain $(1-\epsilon, \delta)$-approximation hardness results for Max-Γ-2Lin$(\mathbb{Z})$ and Max-Γ-3Lin$(\mathbb{Z})$. (Here Γ-3Lin means that all equations are of the form $x_{i_1} + x_{i_2} - x_{i_3} = b$.) Thus we resolve the open question about Max-Γ-2Lin$(\mathbb{Z})$ and give a simpler proof of the Guruswami–Raghavendra [GR07] result. Our results also hold over $\mathbb{R}$ and over "superconstantly large" cyclic groups $\mathbb{Z}_q$ (we are not aware of previously known hardness results over $\mathbb{Z}_q$ when q is superconstant and prime). The results also have an essentially optimal quantitative tradeoff between ϵ, δ, and the magnitudes of the "right-hand side constants" b.

To state our two theorems, let us define B-Bounded-Max-Γ-2Lin and B-Bounded-Max-Γ-3Lin to be the special cases of Max-Γ-3Lin and Max-Γ-2Lin in which all right-hand side constants b are integers satisfying $|b| \le B$. Given an instance $\mathcal{I}$ of Max-Γ-kLin with integer constants b, we use the notation $\mathsf{Opt}_R(\mathcal{I})$ to denote the maximum weight of equations that can be satisfied when the equations are evaluated over R.

Theorem I.1. *For all constant $\epsilon, \gamma, \kappa > 0$ and constant $q \in \mathbb{N}$, given a q-*Bounded-Max-Γ-2Lin *instance $\mathcal{I}$ it is* UG*-hard to distinguish the following two cases:*

- (Completeness.) *There is a $(1-\epsilon-3\gamma)$-good assignment over $\mathbb{Z}$; i.e., $\mathsf{Opt}_{\mathbb{Z}}(\mathcal{I}) \ge 1-\epsilon-3\gamma$.*
- (Soundness.) *There is no $(q\Lambda_{1-\epsilon}(1/q) + \kappa)$-good assignment over $\mathbb{Z}_q$; i.e., $\mathsf{Opt}_{\mathbb{Z}_q}(\mathcal{I}) \le q\Lambda_{1-\epsilon}(1/q) + \kappa$.*

Note that $q\Lambda_{1-\epsilon}(1/q) \approx (1/q)^{\epsilon/(2-\epsilon)}$ is the same soundness proved by KKMO [KKMO07] for Max-Γ-2Lin$(\mathbb{Z}_q)$.

Theorem I.2. *For all constant $\epsilon, \kappa > 0$ and $q \in \mathbb{N}$, given a q-*Bounded-Max-Γ-3Lin *instance $\mathcal{I}$ it is* NP*-hard to distinguish the following two cases:*

- (Completeness.) *There is a $(1-O(\epsilon))$-good assignment over $\mathbb{Z}$, i.e., $\mathsf{Opt}_{\mathbb{Z}}(\mathcal{I}) \ge 1-O(\epsilon)$.*
- (Soundness.) *There is no $(1/q+\kappa)$-good assignment over $\mathbb{Z}_q$; i.e., $\mathsf{Opt}_{\mathbb{Z}_q}(\mathcal{I}) \le 1/q+\kappa$.*

Note that $\mathsf{Opt}_{\mathbb{Z}}(\mathcal{I}) \le \mathsf{Opt}_{\mathbb{Z}_q}(\mathcal{I})$ since we can convert a δ-good assignment over $\mathbb{Z}$ to a δ-good assignment over $\mathbb{Z}_q$ by reducing the integer solution modulo q. Therefore our hardness results are of the strongest "bicriteria" type: even when promised that there is near-perfect solution over $\mathbb{Z}$, it is hard for an algorithm to find a slightly good solution over $\mathbb{Z}_q$. Indeed, by virtue of Lemma A.1 in Appendix A, by losing just a constant factor in the soundness, we can show that it is also hard for an algorithm to find a slightly good solution over any ring $\{\mathbb{R}, \mathbb{Z}, \mathbb{Z}_{q+1}, Z_{q+2}, \ldots\}$ of the algorithm's choosing. Our results subsume and unify all aforementioned results on Max-3Lin$(\mathbb{Z}_q)$, Max-3Lin$(\mathbb{Z})$, and Max-2Lin$(\mathbb{Z}_q)$, and also provide an optimal UG-hardness result for Max-Γ-2Lin$(\mathbb{Z})$.

II. Preliminaries

A. Notations and Definitions

We write $\mathbb{Z}_q$ for the integers modulo q, and we identify the elements with $\{0, 1, \ldots, q-1\} \in \mathbb{Z}$. We sometimes write $\oplus_q$ for addition of integers modulo q

and $+$ for addition over the integers. For two vectors $x, y \in \mathbb{Z}^n$, both $x \oplus_q y$ and $x + y$ are coordinate-wise add. We will also write Δ_q for the set of probability distributions over $\mathbb{Z}_q$. We can identify Δ_q with the standard $(q-1)$-dimensional simplex in $\mathbb{R}^q$. We also identify an element $a \in \mathbb{Z}_q$ with a distribution in Δ_q, namely, the distribution that puts all of its probability mass on a.

Fix $x \in \mathbb{Z}_q^n$, a random variable $\boldsymbol{y}$ is $(1-\epsilon)$-correlated to x, i.e. $\boldsymbol{y} \sim_{1-\epsilon} x$, if $\boldsymbol{y}$ can be get by rerandomizing each coordinate of x independently with probability ϵ.

We recall some standard definitions from the harmonic analysis of boolean functions (see, e.g., [Rag09]). We will be considering functions of the form $f : \mathbb{Z}_q^n \to \mathbb{R}$. The set of all functions $f : \mathbb{Z}_q^n \to \mathbb{R}$ forms an inner product space with inner product

$$\langle f, g \rangle = \mathop{\mathbf{E}}_{\boldsymbol{x} \sim \mathbb{Z}_q^n}[f(\boldsymbol{x}) \cdot g(\boldsymbol{x})],$$

where $\boldsymbol{x} \sim \mathbb{Z}_q^n$ means that $\boldsymbol{x}$ is uniform randomly chosen from $\mathbb{Z}_q^n$. We also write $\|f\|_2 = \sqrt{\langle f, f \rangle}$ as usual.

The following Efron–Stein decomposition theorem is well-known; see [KKMO07].

Theorem II.1. *Any $f : \mathbb{Z}_q^n \to \mathbb{R}$ can be uniquely decomposed as a sum of functions*

$$f(x) = \sum_{S \subseteq [n]} f^S(x),$$

where

- *$f^S(x)$ depends only on $x_S = (x_i, i \in S)$,*
- *for every $S \subseteq [n]$, for every S' such that $S \backslash S' \neq \emptyset$, and for every $y \in \mathbb{Z}_q^n$, it holds that*

$$\mathop{\mathbf{E}}_{\boldsymbol{x} \sim \mathbb{Z}_q^n}[f^S(\boldsymbol{x}) | \boldsymbol{x}_{S'} = y_{S'}] = 0.$$

Definition II.2 (Influences)**.** For functions $f : \mathbb{Z}_q^n \to \mathbb{R}$, define the influence of the i-th coordinate on f to be

$$\mathbf{Inf}_i(f) = \sum_{S \ni i} \|f^S\|_2^2,$$

where $\|f^S\|_2^2 = \mathbf{E}_{\boldsymbol{x}}[f^S(\boldsymbol{x})^2]$. For functions $f : \mathbb{Z}_q^n \to \Delta_m$, let

$$\mathbf{Inf}_i(f) = \sum_{a \in \mathbb{Z}_m} \mathbf{Inf}_i(f_a),$$

where $f_a(x) = f(x)_a, \forall x \in \mathbb{Z}_q^n$.

Definition II.3 (Noise operators)**.** For functions $f : \mathbb{Z}_q^n \to \mathbb{R}$, define the noise operator $T_{1-\eta}$ to be

$$T_{1-\eta} f(x) = \mathop{\mathbf{E}}_{\boldsymbol{y} \sim_{1-\eta} x}[f(\boldsymbol{y})].$$

For functions $f : \mathbb{Z}_q^n \to \Delta_m$, let $T_{1-\eta}$ be the noise operator so that $(T_{1-\eta} f)_a = T_{1-\eta}(f_a), \forall a \in \mathbb{Z}_q$.

Definition II.4 (Noisy-influences)**.** For functions $f : \mathbb{Z}_q^n \to \mathbb{R}$ and functions $f : \mathbb{Z}_q^n \to \Delta_m$, define the $(1-\eta)$-noisy-influence of the i-th coordinate of f to be

$$\mathbf{Inf}_i^{(1-\eta)}(f) = \mathbf{Inf}_i(T_{1-\eta} f).$$

Fact II.5. *For functions $f : \mathbb{Z}_q^n \to \Delta_m$, we have*

$$\sum_{i \in \mathbb{Z}_q} \mathbf{Inf}_i^{(1-\eta)}(f) \leq 1/\eta.$$

Fact II.6. *Let $f^{(1)}, \ldots, f^{(t)}$ be a collection of functions $\mathbb{Z}_q^n \to \mathbb{R}^m$. Then*

$$\mathbf{Inf}_i^{(1-\eta)}\left[\operatorname*{avg}_{k \in [t]}\left\{f^{(k)}\right\}\right] \leq \operatorname*{avg}_{k \in [t]}\left\{\mathbf{Inf}_i^{(1-\eta)}[f^{(k)}]\right\}.$$

Here for any $c_1, c_2, \ldots c_t \in \mathbb{R}$ (or $\mathbb{R}^m$), we use the notation $\operatorname{avg}(c_1, \ldots, c_t)$ to denote their average:

$$\frac{\sum_{i=1}^t c_i}{t}.$$

Definition II.7 (Noise stability)**.** For functions $f : \mathbb{Z}_q^n \to \mathbb{R}$, define its stability against ϵ noise to be

$$\mathbf{Stab}_{1-\epsilon}[f] = \mathop{\mathbf{E}}_{\boldsymbol{x} \sim \mathbb{Z}_q^n, \boldsymbol{y} \sim_{1-\epsilon} \boldsymbol{x}}[f(\boldsymbol{x}) f(\boldsymbol{y})].$$

One tool we need is the *Majority Is Stablest Theorem* from [MOO05]. (We state here a version using a small noisy-influences assumption rather than a small "low-degree influences" assumption; see, e.g., Theorem 3.2 in [Rag09] for a sketch of the small modification to [MOO05] needed.)

Theorem II.8. *For every function $f : \mathbb{Z}_q^n \to [0,1]$ such that $\mathbf{Inf}_i^{(1-\eta)}[f] \leq \tau$ for all $i \in [n]$, Let $\mu = \mathbf{E}[f]$. Then for any $0 < \epsilon < 1$,*

$$\mathbf{Stab}_{1-\epsilon}[f] \leq \Lambda_{1-\epsilon}(\mu) + e(\tau, q, \eta).$$

Here if we fix η, q, $e(\tau, \eta, q)$ goes to 0 when τ goes to 0.

In the above theorem, the quantity $\Lambda_{1-\epsilon}(\mu)$ is defined to be $\mathbf{Pr}[\boldsymbol{x}, \boldsymbol{y} \leq t]$ when $(\boldsymbol{x}, \boldsymbol{y})$ are joint standard Gaussians with covariance $1-\epsilon$ and t is defined by $\mathbf{Pr}[\boldsymbol{x} \leq t] = \mu$.

III. Review of proofs of Max-Γ-2Lin($\mathbb{Z}_q$) and Max-Γ-3Lin($\mathbb{Z}_q$) hardness

As mentioned, we prove Theorems I.1 and I.2 by fairly easy modifications of the known hardness results for Max-Γ-2Lin($\mathbb{Z}_q$) and Max-Γ-3Lin($\mathbb{Z}_q$), due respectively to Khot–Kindler–Mossel–O'Donnell [KKMO07] and Håstad [Hås01]. In this section, we review several places in the two proofs that are related to our modifications. We also assume the reader's familiarity with these works.

A. Max-Γ-2Lin

Let us begin with Max-Γ-2Lin. As shown in [KKMO07], to prove UG-hardness of $(1-\epsilon,\delta)$-approximating Max-Γ-2Lin($\mathbb{Z}_q$) for constant κ and q, where $\delta = q\Lambda_{1-\epsilon}(1/q)+\kappa$, it suffices to construct a "Dictator vs. Small Low-Degree-Influences Test" (or, Dictator Test for short) for functions $f\colon \mathbb{Z}_q^L \to \Delta_q$ which uses Γ-2Lin constraints and has completeness $1-\epsilon$, soundness δ. We recall the definition of Dictator Test as follows.

Generally speaking, a $1-\epsilon$ vs. δ Dictator Test for functions $f : \mathbb{Z}_q^L \to \mathbb{Z}_q$ is defined by a distribution over Γ-2Lin constraints (over the entries of f). We say f passes the test when a random constraint (from the distribution) is satisfied by f. At the completeness side, all the L dictators (i.e., $f(x) = x_i$ for some $i \in L$) pass the test with probability at least $1-\epsilon$. At the soundness side, all functions with small noisy-influences (on all coordinates) pass the test with probability at most δ. KKMO indeed needs to construct a Dictator Test for functions of distributions, i.e., for $f : \mathbb{Z}_q^L \to \Delta_q$, where whenever the test refers an entry $f(x)$ for an element in $\mathbb{Z}_q$, it randomly samples an element from the distribution $f(x)$.

The Dictator Test used by KKMO is indeed a noise-stability test. Intuitively, dictator functions have high noise stability, while functions far from dictators have low noise stability. Note that this intuition is true only for balanced functions, as constant functions are far from dictators but very noise stable. Therefore, KKMO used the "folding" trick (which was introduced in [Hås01]) to ensure that f outputs $1,2,\ldots,q$ with the same probability.

B. Max-Γ-3Lin

Let us move on to Max-Γ-3Lin and our proof of Theorem I.2. Håstad essentially showed that to prove NP-hardness of $(1-\epsilon, 1/q+\kappa)$-approximating Max-Γ-3Lin($\mathbb{Z}_q$) for constant q, it suffices to construct a "Matching-Dictator Test" on two functions for $f\colon \mathbb{Z}_q^K \to \mathbb{Z}_q$, $g\colon \mathbb{Z}_q^L \to \mathbb{Z}_q$ and $\pi\colon L \to K$. The test is defined by a distribution over $\boldsymbol{x} \in \mathbb{Z}_q^K, \boldsymbol{y} \in \mathbb{Z}_q^L, \boldsymbol{z} \in \mathbb{Z}_q^L$ with the check $f(x)+g(y)-g(z) = c \mod q$. Håstad's Test has the following completeness and soundness promises:

- If $f(x) = x_i$ and $g(y) = y_j$ such that $\pi(i) = j$, then f and g passes with probability $1-\epsilon$.
- If f and g passes the test with probability $1/q+\kappa$, then there is a randomized procedure that "decodes" f into a coordinate $i \in L$ and g into a coordinate $j \in K$ such that $\pi(i) = j$ with constant probability depending only on q, ϵ, κ and independent of L, K, π. Also note that the decoding processes for f and g should be independent from each other.

Håstad constructed the following test: choose $\boldsymbol{x} \in \mathbb{Z}_q^K$ and $\boldsymbol{y} \in \mathbb{Z}_q^L$ uniformly and independently, define $\boldsymbol{z} \in \mathbb{Z}_q^L$ to be $\boldsymbol{z} = \boldsymbol{y} \oplus_q (\boldsymbol{x} \circ \pi)$, where $(\boldsymbol{x}\circ\pi)_i := \boldsymbol{x}_{\pi(i)}$, let $\boldsymbol{z}'$ be $(1-\epsilon)$-correlated to $\boldsymbol{z}$, and test $f(\boldsymbol{x}) \oplus_q g(\boldsymbol{y}) = g(\boldsymbol{z}')$. Such a test does not work when $f \equiv 0$; thus Håstad introduced and used his method of folding (which was also used [KKMO07]) to ensure that f outputs $1,2,\ldots,q$ with equal probability.

IV. Overview of our proofs

As mentioned, we obtain Theorems I.1 and I.2 by modifying the KKMO [KKMO07] and Håstad [Hås01] proofs. In this section we describe the idea of the modifications.

A. Active folding

The usual folding trick [Hås01] enforces that f is balanced by replacing references to $f(x_1,\ldots,x_L)$ with references to $f(x_1 \oplus_q x_{j^*}, x_2 \oplus_q x_{j^*}, \ldots, x_L \oplus_q x_{j^*}) \oplus_q (-x_{j^*})$ for some arbitrary $j^* \in [L]$. (I.e., the reduction only uses q^{L-1} variables to represent f as opposed to q^L. Note that this makes the test's constraints of the form $f(x) \oplus_q b = f(x') \oplus_q b'$, but this is still of Γ-2Lin type. We call this trick *static folding*.

Let us explain the alternative to "static folding" which we call *active folding*. Active folding is nothing more than building the folding directly into the test. We feel that this is slightly more natural than static folding, and as we will see it proves to be more flexible. In the KKMO context of Max-Γ-2Lin($\mathbb{Z}_q$), active folding means that the test additionally chooses $\boldsymbol{c}, \boldsymbol{c}' \sim \mathbb{Z}_q$ uniformly and independently, and then it checks the Γ-2Lin constraint

$$f(\boldsymbol{x} \oplus_q (\boldsymbol{c},\ldots\boldsymbol{c})) \oplus_q (-\boldsymbol{c}) = f(\boldsymbol{x}' \oplus_q (\boldsymbol{c}',\ldots\boldsymbol{c}')) \oplus_q (-\boldsymbol{c}')$$

rather than $f(\boldsymbol{x}) = f(\boldsymbol{x}')$. To analyze the KKMO test with active folding, first note that completeness does not change. As for the soundness analysis, given a function $f : \mathbb{Z}_q^L \to \Delta_q$ we introduce $\widetilde{f}\colon \mathbb{Z}_q^L \to \Delta_q$ defined by

$$\widetilde{f}(x)_a = \mathop{\mathbf{E}}_{\boldsymbol{c}\sim\mathbb{Z}_q}[f(x \oplus_q (\boldsymbol{c},\ldots,\boldsymbol{c}))_{a\oplus_q \boldsymbol{c}}]. \qquad (1)$$

Then the probability f satisfies the test with active folding is precisely the probability that $\widetilde{f}$ satisfies the $\widetilde{f}(\boldsymbol{x}) = \widetilde{f}(\boldsymbol{x}')$ test (in the sense of randomized functions), namely $\mathbf{Stab}_{1-\epsilon}[\widetilde{f}]$. We can now proceed with the KKMO analysis; the key is that we still have $\mathbf{E}[\widetilde{f}_a] = 1/q$ for all $a \in \mathbb{Z}_q$. To see this, take $Q = q$ in the following lemma:

Lemma IV.1. *Let $f : \mathbb{Z}_Q^L \to \Delta_q$ and suppose $\widetilde{f} : \mathbb{Z}_Q^L \to \Delta_q$ is defined as in (1). Then $\mathbf{E}[\widetilde{f}_a] = 1/q$ for all $a \in \mathbb{Z}_q$.*

Proof: We have

$$\mathbf{E}[\widetilde{f}_a] = \mathop{\mathbf{E}}_{\boldsymbol{x}\sim\mathbb{Z}_Q^K,\boldsymbol{c}\sim\mathbb{Z}_q}\left[f(\boldsymbol{x}\oplus_Q(\boldsymbol{c},\ldots,\boldsymbol{c}))_{a\oplus_q\boldsymbol{c}}\right].$$

Write $\widetilde{\boldsymbol{x}} = \boldsymbol{x}\oplus_Q(\boldsymbol{c},\ldots,\boldsymbol{c}) \in \mathbb{Z}_Q^K$. The distribution of $\widetilde{\boldsymbol{x}} \mid (\boldsymbol{c}=c)$ is uniform on $\mathbb{Z}_Q^K$ for every c. In other words, $\widetilde{\boldsymbol{x}}$ and $\boldsymbol{c}$ are independent. Thus

$$\begin{aligned}\mathbf{E}[\widetilde{f}_a] &= \mathop{\mathbf{E}}_{\widetilde{\boldsymbol{x}}}\left[\mathop{\mathbf{E}}_{\boldsymbol{c}}\left[f(\widetilde{\boldsymbol{x}})_{a\oplus_q\boldsymbol{c}}\right]\right]\\ &= \mathop{\mathbf{E}}_{\widetilde{\boldsymbol{x}}}\left[(1/q)\sum_{b\in\mathbb{Z}_q}[f(\widetilde{\boldsymbol{x}})_b]\right] = \mathop{\mathbf{E}}_{\widetilde{\boldsymbol{x}}}[1/q] = 1/q.\end{aligned}$$

■

B. Modifying the KKMO proof

We now describe how to obtain Theorem I.1. Let us first ask: Why does the KKMO reduction (with active folding) not prove Theorem I.1 already? The soundness statement of Theorem I.1 would hold since it is over $\mathbb{Z}_q$. The problem is in the completeness statement: a dictator $f:\mathbb{Z}_q^L\to\mathbb{Z}$, $f(x)=x_i$ does not satisfy the the KKMO test with probability close to 1. The reason is that folding may introduce *wrap-around* in $\mathbb{Z}_q$. More specifically (and ignoring the ϵ noise), the KKMO test with active folding will check

$$(\boldsymbol{x}_i+\boldsymbol{c} \bmod q)-\boldsymbol{c} = (\boldsymbol{x}_i+\boldsymbol{c}' \bmod q)-\boldsymbol{c}' \qquad (2)$$

over the integers, and this is only satisfied if both $\boldsymbol{x}_i+\boldsymbol{c}$ and $\boldsymbol{x}_i+\boldsymbol{c}'$ wrap around, or neither does: probability $1/2$. (The situation with static folding is similar.)

Sketch of a first fix: There is a simple way to somewhat fix the completeness: choose $\boldsymbol{c}$ and $\boldsymbol{c}'$ from a range smaller than $\{0,1,\ldots,q-1\}$. E.g., if we choose $\boldsymbol{c}$ and $\boldsymbol{c}'$ independently and uniformly in $\{0,1,\ldots,\lfloor q/t\rfloor\}$, then we get wrap-around in $\boldsymbol{x}_i+\boldsymbol{c}$ with probability at most $1/t$. Hence the dictator $f(x)=x_i$ will satisfy the test (2) over $\mathbb{Z}$ with probability at least $1-2/t$, which we can make close to 1 by taking t large. Now how does this restricted folding affect the soundness analysis? If we redefine the folded function $\widetilde{f}$ appropriately, it is not hard to show that we will have $\mathbf{E}[f_a]\le(t/q)$ for all a. We could then proceed with the KKMO analysis applied to $\widetilde{f}$ and obtain soundness $q\Lambda_{1-\epsilon}(t/q)$. Choosing, say, $t=\log q$ would achieve a good completeness versus soundness tradeoff; roughly $1-\epsilon'$ versus $\widetilde{O}(1/q)^{\epsilon/(2-\epsilon)}$.

A better fix: A slight twist on this idea actually gives the optimal completeness versus soundness tradeoff. Instead of restricting the range of the folding, we simply enlarge the domain of f. Specifically, let $\gamma>0$ be any small constant and define $Q=\lceil q/\gamma\rceil$. To prove Theorem I.1 we run the KKMO reduction with functions f whose domain is $\mathbb{Z}_Q^L$. We still active folding with $\boldsymbol{c}\in\mathbb{Z}_q$. In other words, the test chooses $\boldsymbol{x},\boldsymbol{x}'$ to be $(1-\epsilon)$-correlated strings in $\mathbb{Z}_Q^L$, chooses $\boldsymbol{c},\boldsymbol{c}'\in\mathbb{Z}_q$ uniformly and independently, and outputs the constraint $f(\boldsymbol{x}\oplus_Q(\boldsymbol{c},\ldots,\boldsymbol{c}))-\boldsymbol{c}=f(\boldsymbol{x}'\oplus_Q(\boldsymbol{c}',\ldots,\boldsymbol{c}'))-\boldsymbol{c}'$. Note that this is a q-Bounded-Γ-2Lin constraint. As the "wrap-around" probability is $q/Q\le\gamma$, we have completeness over $\mathbb{Z}$ of at least $1-\epsilon-\gamma$. As for the soundness over $\mathbb{Z}_q$, we now need to consider functions $f:\mathbb{Z}_Q^L\to\Delta_q$. If we introduce the folded function $\widetilde{f}:\mathbb{Z}_Q^L\to\Delta_q$ as in (1), the probability f passes the test over $\mathbb{Z}_q$ is again $\mathbf{Stab}_{1-\epsilon}[\widetilde{f}]$, and we still have $\mathbf{E}[\widetilde{f}_a]=1/q$ by Lemma IV.1. Hence the soundness analysis for Theorem I.1 becomes essentially identical to the soundness analysis for KKMO with active folding. The only tiny difference is that we need to apply the Majority Is Stablest Theorem with domain $\mathbb{Z}_Q^L$ rather than $\mathbb{Z}_q^L$. But Q is still a constant since γ and q are; hence we obtain the claimed $1-\epsilon-\gamma$ completeness over $\mathbb{Z}$ and $q\Lambda_{1-\epsilon}(1/q)$ soundness over $\mathbb{Z}_q$.

C. Modifying the Håstad proof

The modification to Håstad's test needed to obtain Theorem I.2 is similar. If one carries out Håstad's proof using the Efron–Stein decomposition rather than harmonic analysis over $\mathbb{Z}_q$, one sees that the soundness relies entirely on $\mathbf{E}[f_a]=1/q$ for all $a\in\mathbb{Z}_q$. Thus we only need to apply folding to f. Let us examine the Håstad Γ-3Lin test on $f:\mathbb{Z}_q^K\to\mathbb{Z}_q$, $g:\mathbb{Z}_q^L\to\mathbb{Z}_q$, and $\pi:L\to K$. We will use active folding on f, and for simplicity of this discussion ignore the ϵ-noise. The test chooses $\boldsymbol{x}\sim\mathbb{Z}_q^K$ and $\boldsymbol{y}\sim\mathbb{Z}_q^L$ uniformly and independently, defines $\boldsymbol{z}\in\mathbb{Z}_q^L$ by $\boldsymbol{z}=\boldsymbol{y}\oplus_q(\boldsymbol{x}\circ\pi)$ (again, $(\boldsymbol{x}\circ\pi)_i:=\boldsymbol{x}_{\pi(i)}$), chooses $\boldsymbol{c}\sim\mathbb{Z}_q$ uniformly, and finally checks the Γ-3Lin constraint

$$f(\boldsymbol{x}\oplus_q(\boldsymbol{c},\ldots,\boldsymbol{c}))-\boldsymbol{c}+g(\boldsymbol{y})=g(\boldsymbol{z}).$$

Again, if we simply use this reduction in an attempt to prove Theorem I.2, the soundness is fine but the completeness over $\mathbb{Z}$ is a problem due to wrap-around. Indeed, there are two possibilities for wrap-around here: in $\boldsymbol{x}_i+\boldsymbol{c}$ and in $\boldsymbol{y}_j+\boldsymbol{x}_{\pi(j)}$. We mitigate this with the same idea used for Max-Γ-2Lin. Given constants ϵ and q we define constants $Q=\lceil q/\epsilon\rceil$ and $\mathcal{Q}=\lceil Q/\epsilon\rceil$. We enlarge f's domain to $\mathbb{Z}_Q^K$ and g's domain to $\mathbb{Z}_{\mathcal{Q}}^L$. We continue to fold f using $\boldsymbol{c}\sim\mathbb{Z}_q$. Now the two possibilities for wrap-around occur with probability at most ϵ each and hence the completeness over $\mathbb{Z}$ is $1-O(\epsilon)$. Defining $\widetilde{f}:\mathbb{Z}_Q^K\to\Delta_q$ as in (1), we again have $\mathbf{E}[\widetilde{f}_a]=1/q$ for each $a\in\mathbb{Z}_q$ and can carry out the (Efron–Stein-style) Håstad soundness analysis, obtaining soundness $1/q+\kappa$ over $\mathbb{Z}_q$.

Figure 1. **Test $\mathcal{T}$ with parameters ϵ, γ, q for functions on $\mathbb{Z}_Q^K$:**

- Choose $\boldsymbol{x}, \boldsymbol{x}' \sim \mathbb{Z}_Q^K$ to be a pair of $(1-\epsilon)$-correlated random strings.
- Choose $\boldsymbol{c}, \boldsymbol{c}' \sim [q]$ independently and uniformly.
- Define $\widetilde{\boldsymbol{x}} = \boldsymbol{x} \oplus_Q (\boldsymbol{c}, \boldsymbol{c}, \ldots, \boldsymbol{c})$, and define $\widetilde{\boldsymbol{x}}' = \boldsymbol{x}' \oplus_Q (\boldsymbol{c}', \boldsymbol{c}', \ldots, \boldsymbol{c}')$.
- Test the constraint $f(\widetilde{\boldsymbol{x}}) - \boldsymbol{c} = f(\widetilde{\boldsymbol{x}}') - \boldsymbol{c}'$.

V. Dictator Test details

A. Dictator Test for Max-Γ-2Lin

Given constants $\epsilon, \gamma, \kappa > 0$ and $q, K \in \mathbb{Z}^+$, let $Q = \lceil q/\gamma \rceil$. We define the Dictator Test $\mathcal{T}$ for functions f with domain $\mathbb{Z}_Q^K$ as in Figure 1. Let $\mathsf{Val}_{\mathbb{Z}}^{\mathcal{T}}(f)$ be the probability that f passes the test, and let $\mathsf{Val}_{\mathbb{Z}_q}^{\mathcal{T}}(f)$ be the probability that f passes the test over $\mathbb{Z}_q$.

Theorem V.1. *There exists $\tau, \eta > 0$ such that $\mathcal{T}$ is a q-Bounded-Γ-2Lin test with following properties:*

- (Completeness.) *Each of the K dictators $f : \mathbb{Z}_Q^K \to \mathbb{Z}$ has $\mathsf{Val}_{\mathbb{Z}}^{\mathcal{T}}(f) \geq 1 - \epsilon - \gamma$.*
- (Soundness.) *Let $f : \mathbb{Z}_Q^K \to \Delta_q$ and define $\widetilde{f} : \mathbb{Z}_Q^K \to \Delta_q$ as in (1). Suppose that $\mathbf{Inf}_i^{(1-\eta)}[\widetilde{f}] \leq \tau$ for all $i \in [K]$. Then $\mathsf{Val}_{\mathbb{Z}_q}^{\mathcal{T}}(f) \leq q\Lambda_{1-\epsilon}(1/q) + \kappa$, where $\kappa = \kappa(\tau, Q, \eta) > 0$ can be made arbitrarily small by taking $\tau, \eta > 0$ sufficiently small.*

Theorem V.1 together with the following lemma proves Theorem I.1.

Lemma V.2. *Theorem V.1 implies Theorem I.1 .*

Lemma V.2 is implicit from [KKMO07], and is proved in Appendix B1.

Proof of Theorem V.1: For the Completeness case, we need to analyze for a fixed $i \in [K]$ the probability that

$$(\boldsymbol{x}_i \oplus_Q \boldsymbol{c}) - \boldsymbol{c} = (\boldsymbol{x}'_i \oplus_Q \boldsymbol{c}') - \boldsymbol{c}' \tag{3}$$

holds over $\mathbb{Z}$. We have $\boldsymbol{x}_i = \boldsymbol{x}'_i$ except with probability at most ϵ, and $\boldsymbol{x}_i \leq Q - q$ except with probability at most $q/Q \leq \gamma$. When both of these events occur, equation (3) holds. This proves the completeness.

As for the Soundness case, by Lemma IV.1 we have $\mu_a = \mathbf{E}[\widetilde{f}_a] = 1/q$ for each $a \in \mathbb{Z}_q$. By assumption we have $\mathbf{Inf}_i^{(1-\eta)}[\widetilde{f}_a] \leq \mathbf{Inf}_i^{(1-\eta)}[\widetilde{f}] \leq \tau$. Thus from Theorem II.8 we obtain $\mathbf{Stab}_{1-\epsilon}[\widetilde{f}_a] \leq \Lambda_{1-\epsilon}(1/q) + e(\tau, Q, \eta)$ for each a. Summing this over $a \in \mathbb{Z}_q$ yields

$$\mathbf{Stab}_{1-\epsilon}[\widetilde{f}] \leq q\Lambda_{1-\epsilon}(1/q) + q \cdot e(\tau, Q, \eta).$$

The proof is completed by taking $\kappa = q \cdot e(\tau, Q, \eta)$, since $\mathbf{Stab}_{1-\epsilon}[\widetilde{f}] = \mathsf{Val}_{\mathbb{Z}_q}^{\mathcal{T}}(f)$ by unrolling definitions. ∎

B. Matching Dictator Test for Max-Γ-3Lin

Given constants $\epsilon, \kappa > 0$ and $q, L, K \in \mathbb{Z}$, let $Q = \lceil q/\epsilon \rceil$ and $\mathcal{Q} = \lceil Q/\epsilon \rceil$. In Figure 2, we define the Matching Dictator Test $\mathcal{U}$ for function f with domain $\mathbb{Z}_Q^K$, function g with domain $\mathbb{Z}_{\mathcal{Q}}^L$, and projection $\pi : L \to K$. Let $\mathsf{Val}_{\mathbb{Z}}^{\mathcal{U}}(f, g)$ be the probability that f, g pass the test, and let $\mathsf{Val}_{\mathbb{Z}_q}^{\mathcal{U}}(f, g)$ be the probability that f, g pass the test over $\mathbb{Z}_q$.

Theorem V.3. *$\mathcal{U}$ is a q-Bounded-Γ-3Lin test satisfying:*

- (Completeness.) *If $f : \mathbb{Z}_Q^K \to \mathbb{Z}$ and $g : \mathbb{Z}_{\mathcal{Q}}^L \to \mathbb{Z}$ are matching dictators — i.e., $f(x) = x_{\pi(j)}$ and $g(y) = y_j$ for some $j \in [L]$ — then $\mathsf{Val}_{\mathbb{Z}}^{\mathcal{U}}(f, g) \geq 1 - 5\epsilon$.*
- (Soundness.) *Let $f : \mathbb{Z}_Q^K \to \mathbb{Z}_q$, $g : \mathbb{Z}_{\mathcal{Q}}^L \to \mathbb{Z}_q$ and define $\widetilde{f} : \mathbb{Z}_Q^K \to \Delta_q$ as in (1). Suppose that $\mathsf{Val}_{\mathbb{Z}_q}^{\mathcal{U}}(f, g) \geq 1/q + \kappa$, then there is a randomized "decoding procedure" $\mathcal{D}$ which decodes g to a coordinate $\mathcal{D}(g) \in [L]$ and f to a coordinate $\mathcal{D}(f) \in [K]$ such that $\pi(\mathcal{D}(g)) = \mathcal{D}(f)$ with at least a constant probability $\zeta = \zeta(q, \epsilon, \kappa)$ independent of π, L, K.*

Theorem V.3 together with the following lemma proves Theorem I.2.

Lemma V.4. *Theorem V.3 implies Theorem I.2 .*

Lemma V.4 is proved in Appendix B2.

Proof of Theorem V.3: Define $\widetilde{f} : \mathbb{Z}_Q^K \to \Delta_q$ as in (1). For the completeness case, we need to analyze for a fixed $j \in [L]$ the probability that

$$\boldsymbol{x}'_{\pi(j)} - \boldsymbol{c} + \boldsymbol{y}'_j = \boldsymbol{z}'_j \tag{4}$$

holds over $\mathbb{Z}$. Except with probability at most 3ϵ we have all of

$$\boldsymbol{x}'_{\pi(j)} = \widetilde{\boldsymbol{x}}_{\pi(j)} = \boldsymbol{x}_{\pi(j)} \oplus_Q \boldsymbol{c},$$
$$\boldsymbol{y}'_j = \boldsymbol{y}_j, \qquad \boldsymbol{z}'_j = \boldsymbol{z}_j = \boldsymbol{x}_{\pi(j)} \oplus_{\mathcal{Q}} \boldsymbol{y}_j.$$

Except with probability at most $q/Q \leq \epsilon$ we have $\boldsymbol{x}_{\pi(j)} \leq Q - q$, in which case $\boldsymbol{x}_{\pi(j)} \oplus_Q \boldsymbol{c}$ equals $\boldsymbol{x}_{\pi(j)} + \boldsymbol{c}$. Except with probability at most $Q/\mathcal{Q} \leq \epsilon$ we have $\boldsymbol{y}_j \leq \mathcal{Q} - Q$, in which case $\boldsymbol{x}_{\pi(j)} \oplus_{\mathcal{Q}} \boldsymbol{y}_j = \boldsymbol{x}_{\pi(j)} + \boldsymbol{y}_j$. Thus when all five events occur, equation (4) indeed holds over $\mathbb{Z}$.

Figure 2. **Test $\mathcal{U}$ with parameters ϵ, q for f on $\mathbb{Z}_Q^K$, g on $\mathbb{Z}_Q^L$, $\pi : L \to K$:**

- Choose $\boldsymbol{x} \sim \mathbb{Z}_Q^K$, $\boldsymbol{y} \sim \mathbb{Z}_Q^L$ uniformly and independently.
- Define $\boldsymbol{z} \in \mathbb{Z}_Q^L$ by $\boldsymbol{z} = \boldsymbol{y} \oplus_Q (\boldsymbol{x} \circ \pi)$.
- Choose $\boldsymbol{c} \sim \mathbb{Z}_q$ uniformly and define $\widetilde{\boldsymbol{x}} \in \mathbb{Z}_Q^K$ by $\widetilde{\boldsymbol{x}} = \boldsymbol{x} \oplus_Q (\boldsymbol{c}, \boldsymbol{c}, \ldots, \boldsymbol{c})$.
- Let $\boldsymbol{x}' \in \mathbb{Z}_Q^K$ be $(1-\epsilon)$-correlated to $\widetilde{\boldsymbol{x}}$, let $\boldsymbol{y}' \in \mathbb{Z}_Q^L$ be $(1-\epsilon)$-correlated to $\boldsymbol{y}$, and let $\boldsymbol{z}' \in \mathbb{Z}_Q^L$ be $(1-\epsilon)$-correlated to $\boldsymbol{z}$.
- Test the constraint $f(\boldsymbol{x}') - \boldsymbol{c} + g(\boldsymbol{y}') = g(\boldsymbol{z}')$.

As for the soundness case, write $f' = T_{1-\epsilon}\widetilde{f}$ and $g' = T_{1-\epsilon}g$, where we think of g as $g : \mathbb{Z}_Q^L \to \Delta_q$. By unrolling definitions we have

$$\mathsf{Val}^{\mathcal{U}}_{\mathbb{Z}_q}(f,g) = \sum_{a,b\in\mathbb{Z}_q} \mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{y},\boldsymbol{z}}[f'_a(\boldsymbol{x})g'_b(\boldsymbol{y})g'_{a\oplus_q b}(\boldsymbol{z})].$$

Write $\mu_a = \mathbf{E}[f'_a(\boldsymbol{x})]$. Thus $\mu_a = \mathbf{E}[\widetilde{f}_a] = 1/q$, by Lemma IV.1. We conclude that

$$\begin{aligned}\mathsf{Val}^{\mathcal{U}}_{\mathbb{Z}_q}(f,g) &= \sum_{a,b\in\mathbb{Z}_q} \mathbf{E}[(f'_a(\boldsymbol{x}) - \mu_a)g'_b(\boldsymbol{y})g'_{a\oplus_q b}(\boldsymbol{z})] \\ &\quad + (1/q)\sum_{a,b\in\mathbb{Z}_q} \mathbf{E}[g'_b(\boldsymbol{y})g'_{a\oplus_q b}(\boldsymbol{z})].\end{aligned}$$

The second term above is

$$\begin{aligned}&(1/q)\sum_{a,b\in\mathbb{Z}_q} \mathbf{E}[g'_b(\boldsymbol{y})g'_{a\oplus_q b}(\boldsymbol{z})] \\ =\ &(1/q)\,\mathbf{E}\big[(\textstyle\sum_c g'_c(\boldsymbol{y}))\cdot(\sum_c g'_c(\boldsymbol{z}))\big] \\ =\ &(1/q)\,\mathbf{E}[1\cdot 1] = 1/q,\end{aligned}$$

since g' is Δ_q-valued. Thus to complete the proof it remains to show that if

$$\sum_{a,b\in\mathbb{Z}_q} \mathbf{E}[(f'_a(\boldsymbol{x}) - \mu_a)g'_b(\boldsymbol{y})g'_{a\oplus_q b}(\boldsymbol{z})] \qquad (5)$$

is at least $\kappa > 0$ then we can suitably decode $\widetilde{f}$ and g. Let us now apply the Efron–Stein decomposition to f' and g' with respect to the uniform distributions on their domains. Given $S \subseteq [K]$, $T \subseteq [L]$, for simplicity we write

$$\boldsymbol{F}^S_a = {f'}^S_a(\boldsymbol{x}), \quad \boldsymbol{G}^T_b = {g'}^T_b(\boldsymbol{y}), \quad \boldsymbol{H}^T_{a+b} = {g'}^T_{a\oplus_q b}(\boldsymbol{z}).$$

Thus

$$\begin{aligned}&(5)\\ &= \sum_{a,b\in\mathbb{Z}_q} \mathbf{E}\Big[\Big(\sum_{\emptyset\neq S\subseteq[K]} \boldsymbol{F}^S_a\Big)\Big(\sum_{T\subseteq[L]} \boldsymbol{G}^T_b\Big)\Big(\sum_{U\subseteq[L]} \boldsymbol{H}^U_b\Big)\Big] \\ &= \sum_{a,b\in\mathbb{Z}_q}\ \sum_{\substack{\emptyset\neq S\subseteq[K]\\ T,U\subseteq[L]}} \mathbf{E}[\boldsymbol{F}^S_a\boldsymbol{G}^T_b\boldsymbol{H}^U_{a+b}].\end{aligned}$$

Let us simplify the above. We have $\mathbf{E}[\boldsymbol{F}^S_a\boldsymbol{G}^T_b\boldsymbol{H}^U_{a+b}] = \mathbf{E}[\boldsymbol{F}^S_a \cdot \mathbf{E}[\boldsymbol{G}^T_b\boldsymbol{H}^U_{a+b} \mid \boldsymbol{x}]]$. Note that even if we condition on $\boldsymbol{x}$, the marginals on $\boldsymbol{y}$ and $\boldsymbol{z}$ are uniform on $\mathbb{Z}_Q^L$. It follows from the properties of the Efron–Stein decomposition that $\mathbf{E}[\boldsymbol{G}^T_b\boldsymbol{H}^U_{a+b} \mid \boldsymbol{x}]$ is always 0 if $T \neq U$. Thus

$$(5) = \sum_{a,b\in\mathbb{Z}_q}\ \sum_{\substack{\emptyset\neq S\subseteq[K]\\ U\subseteq[L]}} \mathbf{E}[\boldsymbol{F}^S_a\boldsymbol{G}^U_b\boldsymbol{H}^U_{a+b}].$$

Similarly, conditioned on the U-coordinates of $\boldsymbol{y}$ and $\boldsymbol{z}$, the coordinates of $\boldsymbol{x}$ outside $\pi(U)$ are independent and uniform on $\mathbb{Z}_Q$. Hence $\mathbf{E}[\boldsymbol{F}^S_a\boldsymbol{G}^U_b\boldsymbol{H}^U_{a+b}] = 0$ if $S \not\subseteq \pi(U)$. We conclude that

$$\begin{aligned}(5) &= \sum_{a,b\in\mathbb{Z}_q}\ \sum_{\substack{U\neq\emptyset\\ \emptyset\neq S\subseteq\pi(U)}} \mathbf{E}[\boldsymbol{F}^S_a\boldsymbol{G}^U_b\boldsymbol{H}^U_{a+b}] \\ &= \sum_{a,b\in\mathbb{Z}_q}\ \sum_{U\neq\emptyset} \mathbf{E}[\boldsymbol{F}^{\leq\pi(U)}_a\boldsymbol{G}^U_b\boldsymbol{H}^U_{a+b}],\end{aligned}$$

where we defined $\boldsymbol{F}^{\leq\pi(U)}_a = \sum_{\emptyset\neq S\subseteq\pi(U)} \boldsymbol{F}^S_a$. Shifting the sum over a and b to the inside we obtain

$$\begin{aligned}(5) &= \sum_{U\neq\emptyset} \mathbf{E}\Big[\sum_{a,b\in\mathbb{Z}_q} \boldsymbol{F}^{\leq\pi(U)}_a\boldsymbol{G}^U_b\boldsymbol{H}^U_{a+b}\Big] \\ &\leq \sum_{U\neq\emptyset} \mathbf{E}\left[\sqrt{\textstyle\sum_{a,b}(\boldsymbol{F}^{\leq\pi(U)}_a)^2(\boldsymbol{G}^U_b)^2}\sqrt{\textstyle\sum_{a,b}(\boldsymbol{H}^U_{a+b})^2}\right],\end{aligned}$$

having used Cauchy-Schwarz. We can think of, e.g., $(\boldsymbol{G}^U_0, \ldots, \boldsymbol{G}^U_{q-1})$ as a vector in $\mathbb{R}^q$; writing $\|\boldsymbol{G}^U\|$ for the Euclidean length of this vector (and similarly for $\boldsymbol{F}$ and $\boldsymbol{H}$), the right side above is precisely $\sqrt{q}\sum_{U\neq\emptyset}\mathbf{E}\big[\|\boldsymbol{F}^{\leq\pi(U)}\|\cdot\|\boldsymbol{G}^U\|\cdot\|\boldsymbol{H}^U\|\big]$. Thus

$$\begin{aligned}(5) &\leq \sqrt{q}\sum_{U\neq\emptyset} \mathbf{E}\big[\|\boldsymbol{F}^{\leq\pi(U)}\|\cdot\|\boldsymbol{G}^U\|\cdot\|\boldsymbol{H}^U\|\big] \\ &\leq \sqrt{q}\sum_{U\neq\emptyset} \sqrt{\mathbf{E}[\|\boldsymbol{F}^{\leq\pi(U)}\|^2\|\boldsymbol{G}^U\|^2]}\sqrt{\mathbf{E}[\|\boldsymbol{H}^U\|^2]},\end{aligned}$$

using Cauchy-Schwarz again. Now $\boldsymbol{F}^{\leq\pi(U)}$ depends only on $\boldsymbol{x}$ and $\boldsymbol{G}^U$ depends only on $\boldsymbol{y}$; hence they are independent. Further, since $\boldsymbol{y}$ and $\boldsymbol{z}$ have the same distribution (though they are not independent), the same

is true of $\boldsymbol{G}^U$ and $\boldsymbol{H}^U$. Hence

$$(5) \le \sqrt{q} \sum_{U \neq \emptyset} \sqrt{\mathbf{E}[\|\boldsymbol{F}^{\le \pi(U)}\|^2]\, \mathbf{E}[\|\boldsymbol{G}^U\|^2]}$$

$$\le \sqrt{q} \sqrt{\sum_{U \neq \emptyset} \mathbf{E}[\|\boldsymbol{F}^{\le \pi(U)}\|^2]\, \mathbf{E}[\|\boldsymbol{G}^U\|^2]} \sqrt{\sum_{U \neq \emptyset} \mathbf{E}[\|\boldsymbol{G}^U\|^2]},$$

using Cauchy-Schwarz again. By (generalized) Parseval, $\sum_{U \neq \emptyset} \mathbf{E}[\|\boldsymbol{G}^U\|^2] \le \sum_T \mathbf{E}[\|\boldsymbol{G}^U\|^2] = \mathbf{E}[\|\boldsymbol{G}\|^2] \le 1$, since $\boldsymbol{G}$ takes values in Δ_q. Thus we finally conclude

$$(5) \le \sqrt{q} \sqrt{\sum_{U \neq \emptyset} \mathbf{E}[\|\boldsymbol{F}^{\le \pi(U)}\|^2]\, \mathbf{E}[\|\boldsymbol{G}^U\|^2]}$$

$$= \sqrt{q \sum_{U \neq \emptyset} \mathbf{E}[\|(T_{1-\epsilon}\widetilde{f})^{\le \pi(U)}(\boldsymbol{x})\|^2]\, \mathbf{E}[\|(T_{1-\epsilon}g)^U(\boldsymbol{y})\|^2]}.$$

If $\mathsf{Val}^{\mathcal{U}}_{\mathbb{Z}_q}(f,g) \ge 1/q + \kappa$, then we have $\kappa \le (5)$ and therefore

$$\sum_{U \neq \emptyset} \mathbf{E}[\|(T_{1-\epsilon}\widetilde{f})^{\le \pi(U)}(\boldsymbol{x})\|^2]\, \mathbf{E}[\|(T_{1-\epsilon}g)^U(\boldsymbol{y})\|^2] \ge \kappa^2/q. \quad (6)$$

We now define the decoding procedure. It works in a similar way as in Håstad's work [Hås01], as follows. We sample a random set $S \subseteq [K]$ according to distribution $\mathbf{E}[\|\widetilde{f}^S(\boldsymbol{x})\|^2]$, and let $\mathcal{D}(f) \in S$ uniformly (or an arbitrary element of $[K]$ if $S = \emptyset$). We also sample a random set $T \subseteq [L]$ according to distribution $\mathbf{E}[\|g^S(\boldsymbol{y})\|^2]$, and choose $\mathcal{D}(g) \in T$ uniformly (or an arbitrary element of $[L]$ if $T = \emptyset$). We have

$$\mathbf{Pr}[\pi(\mathcal{D}(g)) = \mathcal{D}(f)]$$

$$\ge \sum_{T, \emptyset \neq S \subseteq \pi(T)} \mathbf{E}[\|\widetilde{f}^S(\boldsymbol{x})\|^2]\, \mathbf{E}[\|g^T(\boldsymbol{y})\|^2] \frac{1}{|S|}$$

$$\ge 2\epsilon \sum_{T, \emptyset \neq S \subseteq \pi(T)} \mathbf{E}[\|\widetilde{f}^S(\boldsymbol{x})\|^2]\, \mathbf{E}[\|g^T(\boldsymbol{y})\|^2] (1-\epsilon)^{2|S|},$$

where in the last step we use the fact $1/|S| \ge 2\epsilon(1-\epsilon)^{2|S|}$. Note that $\mathbf{E}[\|\widetilde{f}^S(\boldsymbol{x})\|^2](1-\epsilon)^{2|S|} = \mathbf{E}[\|T_{1-\epsilon}\widetilde{f}^S(\boldsymbol{x})\|^2]$ and $\mathbf{E}[\|g^T(\boldsymbol{y})\|^2] \ge \mathbf{E}[\|(T_{1-\epsilon}g)^T(\boldsymbol{y})\|^2]$, we have

$$\mathbf{Pr}[\pi(\mathcal{D}(g)) = \mathcal{D}(f)]$$

$$\ge 2\epsilon \sum_{T, \emptyset \neq S \subseteq \pi(T)} \mathbf{E}[\|(T_{1-\epsilon}\widetilde{f})^S(\boldsymbol{x})\|^2]\, \mathbf{E}[\|(T_{1-\epsilon}g)^T(\boldsymbol{y})\|^2]$$

$$= 2\epsilon \sum_{T} \mathbf{E}[\|(T_{1-\epsilon}\widetilde{f})^{\le \pi(T)}(\boldsymbol{x})\|^2]\, \mathbf{E}[\|(T_{1-\epsilon}g)^T(\boldsymbol{y})\|^2]$$

$$\ge 2\epsilon\kappa^2/q,$$

where the second last step is by definition and orthogonality of $(T_{1-\epsilon}\widetilde{f})^{S_1}$ and $(T_{1-\epsilon}\widetilde{f})^{S_2}$ $(S_1 \neq S_2)$, and the last step is by (6). ■

References

[ABSS93] S. Arora, L. Babai, J. Stern, and Z. Sweedyk. The hardness of approximate optima in lattices, codes and linear equations. pages 724–733, 1993. I-A

[ALM+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, May 1998. A.5

[AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *Journal of the ACM*, 45(1):70–122, 1998. A.5

[FGKP06] Vitaly Feldman, Parikshit Gopalan, Subhash Khot, and Ashok Ponnuswami. New results for learning noisy parities and halfspaces. In *Proceedings of the 47th Annual IEEE Symposium on Foundations of Computer Science*, pages 349–359, 2006. I-A

[GR06] Venkatesan Guruswami and Prasad Raghavendra. Hardness of learning halfspaces with noise. In *Proc. 47th IEEE Symposium on Foundations of Computer Science*, pages 543–552, 2006. I-A, A

[GR07] Venkatesan Guruswami and Prasad Raghavendra. A 3-query PCP over integers. In *Proc. 39th ACM Symposium on Theory of Computing*, pages 198–206, 2007. I-A, I-C

[GW95] Michel X. Goemans and David P. Williamson. Improved Approximation Algorithms for Maximum Cut and Satisfiability Problems Using Semidefinite Programming. *J. Assoc. Comput. Mach.*, 42:1115–1145, 1995. I-B

[Hås01] Johan Håstad. Some optimal inapproximability results. *Journal of the ACM*, 48(4):798–859, 2001. I-A, III, III-A, IV, IV-A, V-B

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proc. 34th ACM Symposium on Theory of Computing*, pages 767–775, 2002. I-B, A.3

[KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal inapproximability results for Max-Cut and other 2-variable CSPs? *SIAM Journal on Computing*, 37(1):319–357, 2007. I-B, I-C, II-A, III, III-A, III-B, IV, V-A

[MOO05] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *Proceedings of the 46th Annual IEEE Symposium on Foundations of Computer Science*, pages 21–30, 2005. I-B, I-B, II-A

[Rag09] Prasad Raghavendra. *Approximating NP-hard problems: efficient algorithms and their limits*. PhD thesis, University of Washington, 2009. I-B, II-A, II-A

[Raz98] Ran Raz. A parallel repetition theorem. *SIAM Journal on Computing*, 27(3):763–803, June 1998. A.5

APPENDIX

A. Reductions between Max-kLin(R) problems

Lemma A.1. *Given a q-Bounded-Max-Γ-kLin instance and positive integer $m \geq q$:*

- *When $k = 2$, $\mathsf{Opt}_{\mathbb{Z}}(\mathcal{I}), \mathsf{Opt}_{\mathbb{R}}(\mathcal{I}), \mathsf{Opt}_{\mathbb{Z}_m}(\mathcal{I}) \leq 4 \cdot \mathsf{Opt}_{\mathbb{Z}_q}(\mathcal{I})$.*
- *When $k = 3$, $\mathsf{Opt}_{\mathbb{Z}}(\mathcal{I}), \mathsf{Opt}_{\mathbb{R}}(\mathcal{I}), \mathsf{Opt}_{\mathbb{Z}_m}(\mathcal{I}) \leq 8 \cdot \mathsf{Opt}_{\mathbb{Z}_m}(\mathcal{I})$.*

Proof: It is obvious that the $\mathsf{Opt}_{\mathbb{Z}}$ is a lower bound for $\mathsf{Opt}_{\mathbb{Z}_q}$. It suffice then to show how to convert a δ-good assignment over $\mathbb{Z}_m$ and $\mathbb{R}$ to a $\Omega(\delta)$-good assignment over $\mathbb{Z}$.

First we show the conversion from an assignment over $\mathbb{R}$ to $\mathbb{Z}$. For case of $k = 3$, as is noted in [GR06], suppose one has an δ-good real assignment to a system of equations of the form $x_{i_1} - x_{i_2} + x_{i_3} = b$, $b \in \mathbb{Z}$. If one randomly rounds each variable up or down to an integer, every formerly satisfied equation has probability at least $1/8$ of remaining satisfied.[1] Hence there must exist a $\delta/8$-good integer assignment. For the case of $k = 2$, The reduction from Max-Γ-2Lin($\mathbb{Z}$) to Max-Γ-2Lin($\mathbb{R}$) is even easier and incurs no loss: given a δ-good real assignment, simply dropping the fractional parts yields a δ-good integer assignment.

Next we show the conversion from assignment over $\mathbb{Z}_m$ to $\mathbb{Z}_q$. First let us consider the case of $k = 3$. Suppose one has an δ-good assignment $A : x_i \to Z_m$ to a system of equations of the form

$$x_{i_1} - x_{i_2} + x_{i_3} = b \mod m.$$

Then we know that if $A(x_{i_1}) - A(x_{i_2}) + A(x_{i_3}) = b \mod m$. Notice that $|b| \leq q \leq m$, we must have that $A(x_{i_1}) - A(x_{i_2}) + A(x_{i_3}) \in \{b, b-m, b+m, b+2m\}$ when the assignment is evaluated over $\mathbb{Z}$. If we define assignment $A_1(x_i) = A(x_i) - m, A_2(x_i) = A(x_i) + m$ and $A_3(x_i) = A(x_i) + 2m$ form every x_i. Then it is easy to verify that the best assignment among A, A_1, A_2, A_3 will give a $\delta/4$-good assignment. Essentially, every equations over Z_m satisfiable by A must also be satisfiable by one of A, A_1, A_2, A_3 over $\mathbb{Z}$.

As for the case $k = 2$, we know that for a δ-good assignment A over $\mathbb{Z}_m$, we know that if $A(x_{i_1}) - A(x_{i_2}) = b \mod m$, then $A(x_{i_1}) - A(x_{i_2}) = b - m, b + m$ when evaluated over $\mathbb{Z}$. Therefore, we can randomly set $A'(x_i)$ to be $A(x_i) - m/2$ or $A(x_i) + m/2$. Then we know that A' is at least a $\delta/4$-good assignment over $\mathbb{Z}$.

It is not too hard to see that to see above proof technique also works for $m < q$; in particular, a δ-good assignment for q-Bounded-Max-Γ-kLinon $\mathbb{Z}_m$ implies a $\Omega(\frac{\delta}{(2q/m)^k})$-good assignment on Z_q. ■

[1] In the usual case when the hard instances also have "bipartite" structure, it is not hard to make the loss only a factor of 2 rather than 8.

B. From Dictator Tests to hardness of approximation

1) Proof of Lemma V.2: We start by defining Unique Games and the Unique Games Conjecture.

Definition A.2 (Unique Games)**.** A Unique Game $\mathcal{L}(G(U, V, E), \Sigma, \{\pi_e | e \in E\})$ is a constraint satisfaction problem defined as follows. $G(U, V, E)$ is a bipartite graph whose vertices represent variables and edges represent constraints. The goal is to assign to each vertex a label from the set Σ. The constraint on an edge $e = (u, v) \in E$, where $u \in U, v \in V$, is described by a bijection $\pi_e : \Sigma \to \Sigma$. A labeling $\sigma : U \cup V \to \Sigma$ satisfies the constraint on edge $e = (u, v)$ if and only if $\pi_e(\sigma(v)) = \sigma(u)$. Let $\mathsf{Opt}(\mathcal{U})$ denote the maximum fraction of constraints that can be satisfied by any labeling:

$$\mathsf{Opt}(\mathcal{U}) := \max_{L: U \cup V \to \Sigma} \frac{1}{|E|} \cdot |\{e \in E | L \text{ satisfies } e\}|.$$

Conjecture A.3 (Unique Games Conjecture [Kho02])**.** *For every $\gamma, \delta > 0$, there exists a constant $M = M(\gamma, \delta)$, such that given a Unique Game instance $\mathcal{L}(G(U, V, E), \Sigma, \{\pi_e | e \in E\})$ with $|\Sigma| = M$, it is* NP*-hard to distinguish between these two cases :*

- *YES Case:* $\mathsf{Opt}(\mathcal{L}) \geq 1 - \gamma$.
- *NO Case:* $\mathsf{Opt}(\mathcal{L}) \leq \delta$.

By standard reductions, we can assume the bipartite graph $G(U, V, E)$ is left-regular in the conjecture.

Now we are ready to prove Lemma V.2.

Proof of Lemma V.2: Given a Unique Game instance $\mathcal{L}(G(U, V, E), \Sigma, \{\pi_e | e \in E\})$, and a Dictator Test $\mathcal{T}(\epsilon, \gamma, \kappa, q, K = |\Sigma|)$ described in the lemma statement, we build a q-Bounded-Max-Γ-2Lin instance $\mathcal{I}$ as follows. The variable set consists of all the entries of $g_v : [Q]^{\Sigma} \to \mathbb{Z}, \forall v \in V$, which are supposed Q-ary Long Codes of the labels for $v \in V$, where $Q = q/\gamma$ is defined in the Dictator Test. The equations are placed by the following random process, where the probability of a equation being placed corresponds to its weight.

- Pick a random vertex u and two of its random neighbors of $v, v' \in V$, let $\pi = \pi_{(u,v)}$ and $\pi' = \pi_{(u,v')}$.
- Run the Dictator Test $\mathcal{T}$ an imaginary function f defined on $[Q]^{\Sigma}$, suppose $\mathcal{T}$ chooses to test $f(x) - f(y) = b$.
- Place the equation $(g_v \circ \pi)(x) - (g_{v'} \circ \pi')(y) = b$, where $(g \circ \pi)(x) := g(\pi(x))$.

Completeness. Suppose $\mathsf{Opt}(\mathcal{L}) \geq 1-\gamma$, and σ is a labeling function satisfying $1-\gamma$ fraction of the constraints. Let g_v be the Long Code for $\sigma(v)$, i.e. let $g_v(x) = x_{\sigma(v)}$ for each v. According to the random process shown above, we pick a random equation in $\mathcal{I}$. With probability at least $1-2\gamma$, both of the constraints on (u,v) and (u,v') are satisfied by σ. In this case, both $g_v \circ \pi$ and $g_{v'} \circ \pi'$ are the Long Code for $\sigma(u)$, and $g_v \circ \pi(x) - g_{v'} \circ \pi'(y) = b$ is satisfied with probability $1-\epsilon-\gamma$ by the property of $\mathcal{T}$. In all, at least $1-\epsilon-3\gamma$ fraction (of weight) of the equations are satisfied.

Soundness. Suppose there is a set of functions $g_v : [Q]^\Sigma \to \mathbb{Z}_q$ satisfying more than $q\Lambda_{1-\epsilon}(1/q) + \kappa$ fraction (of weight) of the equations over $\mathbb{Z}_q$. Then there are at least $\kappa/2$ fraction of vertices $u \in U$ such that conditioned on u is picked in the first step of the random process shown above, the equation is satisfied over $\mathbb{Z}_q$ with probability more than $q\Lambda_{1-\epsilon}(1/q) + \kappa/2$. We call such u's "good". For each u, we define $f_u : [Q]^\Sigma \to \Delta_q$ to be $f_u = \mathrm{avg}_{v:(u,v)\in E}\{g_v \circ \pi_{(u,v)}\}$. Since the equations generated after picking u are indeed a Dictator Test $\mathcal{T}$ running on f_u, for good u's, we have $\mathsf{Val}^{\mathcal{T}}_{\mathbb{Z}_q}(f_u) > q\Lambda_{1-\epsilon}(1/q) + \kappa/2$. Therefore, for each good u, there exists $i = i_u \in \Sigma$, such that $\mathbf{Inf}_i^{(1-\eta)}[\widetilde{f_u}] > \tau$. Note that

$$\widetilde{f_u} = \operatorname*{avg}_{v:(u,v)\in E} \{\widetilde{g_v \circ \pi_{(u,v)}}\}.$$

By Fact II.6, we have

$$\begin{aligned}\tau < \mathbf{Inf}_i^{(1-\eta)}[\widetilde{f_u}] &= \mathbf{Inf}_i^{(1-\eta)}\left[\operatorname*{avg}_{v:(u,v)\in E}\{\widetilde{g_v \circ \pi_{(u,v)}}\}\right] \\ &\leq \operatorname*{avg}_{v:(u,v)\in E}\left\{\mathbf{Inf}_i^{(1-\eta)}[\widetilde{g_v \circ \pi_{(u,v)}}]\right\}.\end{aligned}$$

Therefore, for at $\tau/2$ fraction of neighbors v of u, there exists $j = \pi_{(u,v)}(i)$, such that $\mathbf{Inf}_j^{(1-\eta)}(\widetilde{g_v}) > \tau/2$.

Let $\sigma(u) = i_u$ if u is good. For each $v \in V$, let $\mathrm{Cand}(v) = \{i : \mathbf{Inf}_i^{(1-\eta)}(\widetilde{g_v}) > \tau/2\}$. By Fact II.5, we have $|\mathrm{Cand}(v)| < 1/(\tau\eta)$. If $\mathrm{Cand}(v) \neq \emptyset$, let $\sigma(v)$ be a random element in $\mathrm{Cand}(v)$. Now for a good u, there are $\tau/2$ fraction of neighbors v of u such that $j = \pi_{(u,v)}(\sigma(u)) \in \mathrm{Cand}(v)$, therefore the edge (u,v) is satisfied with probability $1/|\mathrm{Cand}(v)| > \tau\eta$. It follows that σ satisfies more than $(\kappa/2)(\tau/2)\tau\eta = \kappa\eta\tau^2/2$ fraction of the constraints in expectation. Therefore there is a labeling satisfying more than $\delta' = \kappa\eta\tau^2/2$ fraction of the constraints. ∎

2) Proof of Lemma V.4: We start by defining Label Cover Games and introducing its hardness.

Definition A.4 (Label Cover Games)**.** A Label Cover Game $\mathcal{C}(G(U,V,E),[K],[L],\{\pi_e|e \in E\})$ is a constraint satisfaction problem defined as follows. $G(U,V,E)$ is a bipartite graph whose vertices represent variables and edges represent the constraints. The goal is to assign to each vertex in U a label from the set $[K]$ and to each vertex in V a label from the set $[L]$. The constraint on an edge $e = (u,v) \in E$ is described by a "projection" $\pi_e : [L] \to [K]$. The projection is onto. A labeling $\sigma : U \to [K], \sigma : V \to [L]$ satisfies the constraint on edge $e = (u,v)$ if and only if $\pi_e(\sigma(v)) = \sigma(u)$. Let $\mathsf{Opt}(\mathcal{C})$ denote the maximum fraction of constraints that can be satisfied by any labeling :

$$\mathsf{Opt}(\mathcal{C}) := \max_{\substack{\sigma:U\to[K]\\ \sigma:V\to[L]}} \frac{1}{|E|} \cdot |\{e \in E | L \text{ satisfies } e\}|.$$

Theorem A.5 (PCP Theorem + Raz's Parallel Repetition Theorem [AS98], [ALM+98], [Raz98])**.** *There exists an absolute constant c such that for every $\delta > 0$, $\mathcal{C}(G(U,V,E),[K],[L],\{\pi_e|e \in E\})$, $K = (1/\delta)^C$, it is* NP*-hard to distinguish between:*

- *YES Case:* $\mathsf{Opt}(\mathcal{C}) = 1$.
- *NO Case:* $\mathsf{Opt}(\mathcal{C}) = \delta$.

Now we are ready to prove Lemma V.4.

Proof of Lemma V.4: Given a Label Cover Game instance $\mathcal{C}(G(U,V,E),[K],[L],\{\pi_e|e \in E\})$, and a Matching Dictator Test $\mathcal{U}(\epsilon,\kappa,q,L,K)$ described in the lemma statement, we build a q-Bounded-Max-Γ-3Lin instance $\mathcal{I}$ as follows. The variable set consists of all the entries of $f_u : [Q]^L \to \mathbb{Z}$ and $g_v : [\mathcal{Q}]^K \to \mathbb{Z}$ for all $u \in U, v \in V$. The equations are the gathering of the Matching Dictator Tests $\mathcal{U}$ for f_u, g_v with projection $\pi_{(u,v)}$ for all $(u,v) \in E$. The weights of the equations are normalised by a factor $1/|E|$.

Completeness. Suppose $\mathsf{Opt}(\mathcal{C}) = 1$, and σ is a labeling function satisfying all the constraints. For all $u \in U, v \in V$, let f_u and g_v be the Long Codes for $\sigma(u), \sigma(v)$ respectively, i.e. let $f_u(x) = x_{\sigma(u)}, g_v(y) = y_{\sigma(v)}$. For each edge $(u,v) \in E$, the Matching Dictator Test $\mathcal{U}$ passes with probability at least $1-\epsilon$. Therefore, at least $1-\epsilon$ fraction (of weight) of the equations are

satisfied.

Soundness. Suppose there is a set of functions $f_u : [Q]^K \to \mathbb{Z}_q, g_v : [\mathcal{Q}]^L \to \mathbb{Z}_q$ satisfying more than $1/q + \kappa$ fraction (of weight) of the equations over $\mathbb{Z}_q$. By averaging argument, for at least $\kappa/2$ fraction of the edges, the corresponding Matching Dictator Test passes with probability more than $1/q + \kappa/2$. Call these edges "good edges". For all $u \in U, v \in V$, let $\sigma(u) = \mathcal{D}(f_u), \sigma(v) = \mathcal{D}(g_v)$. For good edges $e \in E$, the probability that e is satisfied by σ is at least $\zeta = \zeta(q, \epsilon, \kappa)$. It follows that σ satisfies more than $\zeta\kappa/2$ fraction of the constraints in expectation. Therefore there is a labeling satisfying more than $\delta' = \zeta\kappa/2$ fraction of the constraints.

■