Unique Games Over Integers

Ryan O'Donnell* Yi Wu Yuan Zhou
Computer Science Department
Carnegie Mellon University
{odonnell,yiwu,yuanzhou}@cs.cmu.edu

Abstract

Consider systems of two-variable linear equations of the form $x_i - x_j = c_{ij}$, where the c_{ij}'s are integer constants. We show that even if there is an integer solution satisfying at least a $(1-\epsilon)$-fraction of the equations, it is Unique-Games-hard to find an integer (or even real) solution satisfying at least an ϵ-fraction of the equations. Indeed, we show it is Unique-Games-hard even to find an ϵ-good solution modulo *any* integer $m \geq m_0(\epsilon)$ of the algorithm's choosing.

*Supported by NSF grants CCF-0747250 and CCF-0915893, BSF grant 2008477, and Sloan and Okawa fellowships.

1 Introduction

1.1 The equivalence between Unique-Games and Max-2Lin

The Unique Games Conjecture ("UGC") of Khot [Kho02] has been enormously influential in the study of optimization problems, leading to numerous (conditional) optimal inapproximability results (e.g. [DMR06, KR08, GMR08]), especially for constraint satisfaction problems [KKMO07, Rag09]. To recall the conjecture, we need to make several definitions:

Definition 1.1. For $L \in \mathbb{N}$, a $\mathsf{Unique\text{-}Games}_L$ instance consists of a bipartite graph having vertex sets U, V and edge set E, together with a *bijective constraint* $\pi_{uv} : [L] \to [L]$ for each $(u, v) \in E$. In addition, each edge $e \in E$ has a nonnegative *weight* p_{uv}, with $\sum_{(u,v)\in E} p_{uv} = 1$. The algorithmic task is to find an *assignment* $A : (U \cup V) \to [L]$ such that the total weight of satisfied constraints is as large as possible. Here we say that A *satisfies* the constraint π_{uv} if $\pi_{uv}(A(u)) = A(v)$.

Definition 1.2. Given an instance $\mathcal{I}$ of a weighted constraint satisfaction problem (such as $\mathsf{Unique\text{-}Games}_L$) and an assignment A, we define $\mathrm{Val}_{\mathcal{I}}(A)$ to be the total weight of the constraints which A satisfies. We also say that A is *γ-good* if $\mathrm{Val}_{\mathcal{I}}(A) \geq \gamma$. Finally, we define $\mathrm{Opt}(\mathcal{I})$ to be the maximum of $\mathrm{Val}_{\mathcal{I}}(A)$ over all A.

Definition 1.3. For $0 \leq s \leq c \leq 1$, the $\mathsf{Unique\text{-}Games}_L(c, s)$ problem is to distinguish instances $\mathcal{I}$ with $\mathrm{Opt}(\mathcal{I}) \geq c$ from instances $\mathcal{I}$ with $\mathrm{Opt}(\mathcal{I}) < s$.

Definition 1.4. Khot's *Unique Games Conjecture* is that for all small constants $\delta, \epsilon > 0$, there exists a large constant L such that $\mathsf{Unique\text{-}Games}_L(1 - \epsilon, \delta)$ is NP-hard.

Although the Unique-Games problem has a rather abstract definition, Khot, Kindler, Mossel, and O'Donnell [KKMO07] showed that the UGC is equivalent to the hardness of a much more concrete problem: solving two-variable linear equations modulo a large constant.

Definition 1.5. A Max-2Lin instance consists of a set of constraints of the form $v_i - v_j = c_{ij}$. Here $v_1, \ldots, v_n$ are variables and the c_{ij}'s are integer constants. There is also a nonnegative weight for each constraint, with the weights summing to 1.

Definition 1.6. Let $q \in \mathbb{N}$. In the $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_q}$ problem, all equations are interpreted modulo q. We then define $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_q}(c, s)$ to be the task of distinguishing $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_q}$ instances $\mathcal{I}$ with $\mathrm{Opt}(\mathcal{I}) \geq c$ from instances with $\mathrm{Opt}(\mathcal{I}) < s$.

The $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_q}(c, s)$ problem is a special case of the $\mathsf{Unique\text{-}Games}_q(c, s)$ problem. Hence if for all constant $\epsilon, \delta > 0$ there exists a constant q such that $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_q}(1 - \epsilon, \delta)$ is NP-hard, then the UGC is true. The result of [KKMO07] gives an efficient reduction in the opposite direction:

Theorem 1.7. *([KKMO07].) Assuming the UGC, for all $\epsilon, \delta > 0$ there exists q such that $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_q}(1 - \epsilon, \delta)$ is NP-hard.*

An obvious question left open by the [KKMO07] result is whether the UGC also implies hardness of solving two-variable linear equations over the *integers*, rather than over the integers modulo a large constant. I.e., making the obvious definition for $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}}(c, s)$:

Question 1.8. *Is it true that for all constant $\epsilon, \delta > 0$, the $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}}(1-\epsilon, \delta)$ problem is NP-hard assuming the UGC?*

We believe that lack of an additional quantifier over q here gives this question a certain aesthetic appeal.

1.2 Related work

The version of Question 1.8 for Max-3Lin (i.e., equations of the form $v_i - v_j + v_k = c_{ijk}$) took a relatively long time to be resolved. Håstad proved his celebrated NP-hardness result for $\mathsf{Max\text{-}3Lin}_{\mathbb{Z}_q}(1 - \epsilon, 1/q + \delta)$ in 1997 [Hås97]; however, it was not until a decade later that

Guruswami and Raghavendra [GR07] showed that indeed $\mathsf{Max\text{-}3Lin}_{\mathbb{Z}}(1-\epsilon,\delta)$ is NP-hard for all constant $\epsilon, \delta > 0$. A relatively simple observation allowed Guruswami and Raghavendra to also deduce that $\mathsf{Max\text{-}3Lin}_{\mathbb{R}}(1-\epsilon,\delta)$ is NP-hard; here the equations are still of the form $v_i - v_j + v_k = c_{ijk}$ for $c_{ijk} \in \mathbb{Z}$, but the variables can be assigned values in $\mathbb{R}$.

A version of the $\mathsf{Max\text{-}3Lin}_{\mathbb{R}}$ problem is also being studied by Khot and Moshkovitz in ongoing work. In their formulation, called $\mathsf{Robust\text{-}Max\text{-}3Lin}_{\mathbb{R}}$, the constants c_{ijk} are all 0; however certain conditions are placed on how the variables v_i may be assigned real values, so as to eliminate the trivial solution $v_i \equiv 0$. Assuming the UGC, Khot and Moshkovitz [KM10] show that given a system with a $(1-\epsilon)$-good solution, roughly speaking it is NP-hard to find a solution in which a constant fraction of the equations are satisfied to within $\pm\Omega(\sqrt{\epsilon})$. Very recently they have eliminated the need for the UGC. The motivation for their work is the hope of establishing the same sort of result for $\mathsf{Robust\text{-}Max\text{-}2Lin}_{\mathbb{R}}$, a problem closely connected with Unique Games.

1.3 Statement of our result

In this work we show a positive answer to Question 1.8. In fact, our main theorem is the following stronger result:

Theorem 1.9. *Assume the UGC. For any small constants $\epsilon, \delta > 0$, there exists a constant $q = q(\epsilon, \delta) \in \mathbb{N}$ such that the following holds: Given an instance $\mathcal{I}$ of* Max-2Lin *in which the integer constants c_{ij} are in the range $[-q, q]$, it is* NP*-hard to distinguish the following two cases:*

- *There is a $(1-\epsilon)$-good integer assignment to the variables.*
- *There is no assignment to the variables which is δ-good* modulo any integer $m \geq q$.

Assuming $\epsilon, \delta < .1$, it suffices for $q(\epsilon, \delta)$ to be large enough that $\widetilde{O}(1/q)^{\epsilon/(2-\epsilon)} \leq \delta$.

An interesting and somewhat novel aspect of this result is that it gives hardness even for a "multi-objective" problem. In the search version of Theorem 1.9's algorithmic task, although the algorithm is promised there is an extremely good *integer* solution to the given equations, it may attempt to find a slightly good solution modulo *any* $m \geq \widetilde{O}(1/q)^{\epsilon/(2-\epsilon)}$ *of its choosing.* We show that even still, the task is hard assuming the UGC.

From our main result Theorem 1.9, we immediately deduce the following corollaries:

Corollary 1.10. *Assuming the UGC, for all $\epsilon, \delta > 0$ the $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}}(1-\epsilon,\delta)$ problem is* NP*-hard.*

Proof. If there is a δ-good integer assignment to the variables, then this assignment is also δ-good modulo q (or any other integer $m \geq q$). □

Corollary 1.11. *Assuming the UGC, for all $\epsilon, \delta > 0$ there exists q such that the $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_m}(1-\epsilon,\delta)$ problem is* NP*-hard for any $m \geq q$, even for $m = m(n)$ which is super-constant. In particular, the algorithmic task in Theorem 1.9 is equivalent to the UGC.*

Proof. If there is a $(1-\epsilon)$-good integer assignment to the variables, it is also $(1-\epsilon)$-good modulo m. □

Corollary 1.12. *Assuming the UGC, for all $\epsilon, \delta > 0$ the $\mathsf{Max\text{-}2Lin}_{\mathbb{R}}(1-\epsilon,\delta)$ problem is* NP*-hard.*

Proof. Certainly any $(1-\epsilon)$-good integer assignment to the variables is also a $(1-\epsilon)$-good real assignment. Further, as each constraint in Theorem 1.9 is of the form $v_i - v_j = c_{ij} \in \mathbb{Z}$, any δ-good real assignment to the variables v_i can be converted into a δ-good integer assignment simply by dropping all the fractional parts. □

2 Overview of our proof

We now describe the new ideas we introduce to prove Theorem 1.9. In this section, we assume the reader is closely familiar with the proof of the Khot–Kindler–Mossel–O'Donnell (KKMO) UGC-hardness result for $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_q}(1-\epsilon, \delta)$. Our discussions will also not be completely formal.

As KKMO showed, given $\epsilon > 0$ it is sufficient to construct a Dictator Test for functions $f : \mathbb{Z}_q^L \to \mathbb{Z}_q$ using $\mathsf{2Lin}$-constraints, with the following two properties: (i) dictator functions $f(x) = x_i$ pass the test with probability at least $1 - \epsilon$; (ii) any $f : \mathbb{Z}_q^L \to \Delta_q$ with all influences smaller than τ passes the test with probability at most $1/q^{\epsilon/(2-\epsilon)} + \kappa$, where the "error term" $\kappa = \kappa(q, \epsilon, \tau)$ can be made arbitrarily small by taking $\tau > 0$ to be a sufficiently small constant *independent of* L. Here Δ_q is the convexification of $\mathbb{Z}_q$; i.e., the set of all probability distributions over $\mathbb{Z}_q$.

As a first step one might try extending the KKMO analysis to $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_m}$, where m is "super-constant". The essential difficulty is that applying the key tool, the Majority Is Stablest Theorem, to τ-small-influence functions $f : [m]^L \to [0, 1]$ introduces an error term $\kappa(m, \epsilon, \tau)$ which depends on m. If m is super-constant, even as a function of L, this will cause the KKMO reduction from $\mathsf{Unique\text{-}Games}_L$ to fail; in particular, it means that in the soundness case, one would decode such f's to $\omega_L(1)$ many labels in $[L]$, which is unacceptable.

Since we presumably must use the Majority Is Stablest Theorem, and since we also care about constraints modulo a super-constant m, we are led to consider Dictator Tests for functions $f : [q]^L \to \mathbb{Z}_m$. We are not aware of any prior work on testing such functions, with differing domain and range (arguably, the work on hardness of ordering constraints [GMR08] has some of the same flavor). An initial difficulty in working with such functions is that the usual method of "folding" no longer makes sense. Our first observation is that one need not fold by the usual method of restricting the domain by a factor of q; instead, one can build folding directly into the KKMO test. I.e., KKMO's result could be obtained via the following Dictator Test for functions $f : \mathbb{Z}_q^L \to \mathbb{Z}_q$: Choose $\boldsymbol{x}, \boldsymbol{x}' \sim \mathbb{Z}_q^L$ to be $(1 - \epsilon)$-correlated random strings, choose also $\boldsymbol{c}, \boldsymbol{c}' \in \mathbb{Z}_q$ uniformly and independently, and then test the $\mathsf{2Lin}$ constraint

$$f(\boldsymbol{x} + (\boldsymbol{c}, \boldsymbol{c}, \ldots, \boldsymbol{c})) - \boldsymbol{c} = f(\boldsymbol{x}' + (\boldsymbol{c}', \boldsymbol{c}', \ldots, \boldsymbol{c}')) - \boldsymbol{c}'. \tag{1}$$

To analyze the soundness of this test, one introduces the "randomized (or averaged) function" $g : \mathbb{Z}_q^L \to \Delta_q$ defined by $g(x) = g(x + (\boldsymbol{c}, \ldots, \boldsymbol{c})) - \boldsymbol{c}$, in which case the probability that f passes the test is $\mathbb{S}_{1-\epsilon}[g]$. One then observes that $\mathbf{E}[g_a(\boldsymbol{x})] = 1/q$ for each coordinate output function $g_a : \mathbb{Z}_q^L \to [0, 1]$, $a \in \mathbb{Z}_q$. Thus one can apply the Majority Is Stablest to bound $\mathbb{S}_{1-\epsilon}[g]$ by

$$q(\Gamma(1/q) + \kappa(q, \epsilon, \tau)) \leq (1/q)^{\epsilon/(2-\epsilon)} + o_L(1),$$

as necessary.

We will show how to extend this analysis to functions $f : [q]^L \to \mathbb{Z}_m$, where $m \geq q$. Proceeding with the same "built-in folding", we obtain the function $g : [q]^L \to \Delta_m$ which has the property that $\mathbf{E}[g_a(\boldsymbol{x})] \leq 1/q$ for each $a \in [m]$. Our main technical result, Lemma 4.2, shows that this is sufficient to prove

$$\mathbb{S}_{1-\epsilon}[g] = \sum_{a \in [m]} \mathbb{S}_{1-\epsilon}[g_a] \leq (1/q)^{\epsilon/(2-\epsilon)} + q^{\log q} \kappa(q, \epsilon, \tau) = (1/q)^{\epsilon/(2-\epsilon)} + o_L(1).$$

The key point here is that the error term does not depend at all on m, and hence the overall analysis works even for m super-constant. To evade dependence on m, the idea is that one can obtain the bound $\mathbb{S}_{1-\epsilon}[g_a] \leq \mathbf{E}[g_a](1/q)^{\epsilon/2}$ without any small-influences assumption at all if $\mathbf{E}[g_a] \leq q^{-\log q}$; one only needs to use hypercontractivity.

These ideas let us obtain UG-hardness of $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_m}(1 - \epsilon, \delta)$ even for super-constant m. To complete the proof of our main Theorem 1.9, we need to improve the completeness aspect of the Dictator Test so that even integer-valued dictators $f : [q]^L \to \mathbb{Z}$ pass with probability close to 1. An observation here is that an integer-valued dictator $f(x) = x_i$ already passes our test with probability close to $1/2$: Ignoring the ϵ-noise, the test (1) fails only if one of $\boldsymbol{x}_i + \boldsymbol{c}$, $\boldsymbol{x} + \boldsymbol{c}'$ "wraps around" modulo q but the other doesn't.

There is a very simple idea for decreasing the probability of such wraparound: choose $\boldsymbol{c}$ and $\boldsymbol{c}'$ from a range smaller than $[q]$. E.g., if we choose $\boldsymbol{c}, \boldsymbol{c}' \sim [q/t]$, then we get wrap-around in $\boldsymbol{x}_i + \boldsymbol{c}$ with probability at most $1/t$. Hence integer-valued dictators $f : [q]^L \to \mathbb{Z}$, $f(x) = x_i$ will pass the test in (1) with probability at least $1 - \epsilon - 2/t$. How does this restricted folding affect the soundness analysis? It means that the associated randomized function $g : [q]^L \to \Delta_m$ will only satisfy $\mathbf{E}[g_a] \leq t/q$ for each $a \in [m]$, rather than $\mathbf{E}[g_a] \leq 1/q$. But this is still sufficient for our technical Lemma 4.2 to bound $\mathbb{S}_{1-\epsilon}[g]$ by roughly $(t/q)^{\epsilon/(2-\epsilon)}$. Thus by taking $t = \log(q)$, say, we get a 2Lin-based Dictator Test having integer-valued completeness $1 - \epsilon - O(1/\log(q))$ and $\mathbb{Z}_m$-valued soundness $\widetilde{O}(1/q)^{\epsilon/(2-\epsilon)}$ for any $m \geq q$. This suffices to establish our main Theorem 1.9.

2.1 Comparison with Guruswami–Raghavendra

Here we briefly compare our methods with those Guruswami and Raghavendra [GR07] used to establish hardness for $\mathsf{Max\text{-}3Lin}_{\mathbb{Z}}$. Although they also mentioned $\mathsf{Max\text{-}3Lin}_{\mathbb{Z}_m}$ for very large m in the overview of their work, their methods are somewhat more integer-specific than ours. In particular, they worked with Dictator Tests on functions $f : \mathbb{Z}_+^L \to \mathbb{Z}$, using a certain exponential distribution on the domain $\mathbb{Z}_+$. (Ultimately, of course, they truncated the distribution to a finite range.) This necessitated introducing and analyzing a somewhat technical method of decoding functions f to coordinates associated to sparse Fourier frequencies $\omega \in [0, 2\pi]^L$ with large Fourier coefficients.

Guruswami and Raghavendra also described their Dictator Tests as "derandomized versions" of Håstad's tests, where the amount of randomness of the test depends only on the soundness. The same could be said of our result vis-a-vis KKMO's Dictator Tests: we get $\mathsf{Max\text{-}2Lin}_{\mathbb{Z}_m}$ Dictator Tests in which the size of the domain elements, q, depends only on the desired soundness of the test.

2.2 Outline for the rest of the paper

In Section 3 we will introduce the notation and technical tools we will need in the remainder of the paper. The main body of our work is in Section 4 in which we first prove our technical Lemma 4.2, and then we state and analyze our new Dictator Test of $\mathsf{Max\text{-}2Lin}$ type. Although it is essentially known that the construction of this Dictator Test is sufficient to prove Theorem 1.9, we give a proof of this deduction in Appendix A for completeness.

3 Definitions and analytic tools

3.1 Notation

For $r \in \mathbb{R}^+$ we let $[r]$ denote $\{1, 2, \ldots, \lfloor r \rfloor\}$. Given $m \in \mathbb{N}$ we write $\oplus_m$ for addition modulo m. It will also be convenient to use the following slightly unusual notation:

Definition 3.1. We write $\mathbb{Z}_m$ for the group of integers modulo m. We will also sometimes identify this set with $[m] \subset \mathbb{Z}$, *not* with the more standard $\{0, 1, \ldots, m-1\}$. Finally, we extend the notation to $m = \infty$, in which case we understand $\mathbb{Z}_m$ to mean simply the integers, $\mathbb{Z}$.

Definition 3.2. We write Δ_m for the set of probability distributions over $\mathbb{Z}_m$ with finite support; when $m \neq \infty$ we can identify Δ_m with the standard $(m-1)$-dimensional simplex in $\mathbb{R}^m$. We also identify an element $a \in \mathbb{Z}_m$ with a distribution in Δ_m, namely, the distribution that puts all of its probability mass on a.

We write all random variables in **boldface**. If $x \in \Sigma^n$ where Σ is a finite alphabet, we say that the random string $\boldsymbol{y} \in \Sigma^n$ is *ρ-correlated* to x if it is chosen as follows: $\boldsymbol{y}_i$ is set to

x_i with probability ρ and is set to a uniformly random element of Σ with probability $1-\rho$, independently for each $1 \le i \le n$. If $\boldsymbol{x}$ is a uniformly random string in Σ^n and $\boldsymbol{y}$ is ρ-correlated to $\boldsymbol{x}$, we say that $\boldsymbol{x}, \boldsymbol{y}$ are a ρ-correlated pair of random strings.

By $\log(t)$, we will always mean the natural logarithm of t.

3.2 Noise stability and influences

We now recall some standard definitions from the analysis of boolean functions (see, e.g., [Rag09]). We will be considering functions of the form $f : [q]^n \to \mathbb{R}^m$, where $q, n, m \in \mathbb{N}$. We will also allow $m = \infty$, in which case we interpret the range as all sequences in $\mathbb{R}^{\mathbb{Z}}$ with at most finitely many nonzero coordinates. The set of all functions $f : [q]^n \to \mathbb{R}^m$ forms an inner product space with inner product

$$\langle f, g \rangle = \mathop{\mathbf{E}}_{\boldsymbol{x} \sim [q]^n} [\langle f(\boldsymbol{x}), g(\boldsymbol{x}) \rangle];$$

here we mean that $\boldsymbol{x}$ is uniformly random and the $\langle \cdot, \cdot \rangle$ inside the expectation is the usual inner product in $\mathbb{R}^m$. We also write $\|f\| = \sqrt{\langle f, f \rangle}$ as usual.

For $0 \le \rho \le 1$, we define T_ρ to be the linear operator on this inner product space given by

$$T_\rho f(x) = \mathop{\mathbf{E}}_{\boldsymbol{y}}[f(\boldsymbol{y})],$$

where $\boldsymbol{y}$ is a random string in $[q]^L$ which is ρ-correlated to x. We define the *noise stability of f at ρ* to be

$$\mathbb{S}_\rho[f] = \langle f, T_\rho f \rangle.$$

For $i \in [n]$, we define the *influence of i on* $f : [q]^n \to \mathbb{R}^m$ to be

$$\mathrm{Inf}_i[f] = \mathop{\mathbf{E}}_{\boldsymbol{x}_1, \dots, \boldsymbol{x}_{i-1}, \boldsymbol{x}_{i+1}, \dots, \boldsymbol{x}_n \sim [q]} \left[\mathop{\mathbf{Var}}_{\boldsymbol{x}_i \sim [q]} [f(\boldsymbol{x})] \right],$$

where $\mathbf{Var}[f]$ is defined to be $\mathbf{E}[\|f\|^2] - \|\mathbf{E}[f]\|^2$. More generally, for $0 \le \eta \le 1$ we define the *η-noisy-influence of i on f* to be

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

One may observe that

$$\mathrm{Inf}_i^{(1-\eta)}[f] = \sum_{j=1}^{m} \mathrm{Inf}_i^{(1-\eta)}[f_j],$$

where $f_j : [q]^n \to \mathbb{R}$ denotes the jth-coordinate output function of f. (When $m = \infty$ the sum should be over $j \in \mathbb{Z}$.)

We will need the following easy "convexity of noisy-influences" fact:

Proposition 3.3. *Let $f^{(1)}, \dots, f^{(t)}$ be a collection of functions $[q]^n \to \mathbb{R}^m$. Then*

$$\mathrm{Inf}_i^{(1-\eta)} \left[\mathop{\mathrm{avg}}_{k \in [t]} \left\{ f^{(k)} \right\} \right] \le \mathop{\mathrm{avg}}_{k \in [t]} \left\{ \mathrm{Inf}_i^{(1-\eta)}[f^{(k)}] \right\}.$$

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

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

Since $T_{1-\eta}$ is a linear operator, it suffices to prove this for $\eta = 0$ (i.e., for non-noisy influences). This is essentially done within the proof of Theorem 11 in [KKMO07]; see also [Rag09, Proposition 3.0.13].

3.3 Hypercontractivity and Majority Is Stablest

To bound the noise stability of functions we will use two tools. The first is the *hypercontractivity* of the T_ρ operator acting on functions $[q]^L \to \mathbb{R}$. The optimal hypercontractivity bound in this setting was determined by Diaconis and Saloff-Coste [DSC96] (via their Theorems 3.5.ii and A.1, along with Hölder duality):

Theorem 3.4. *Let $q \geq 2$, $f : [q]^n \to \mathbb{R}$, and $0 \leq \epsilon < 1$. Then*

$$\|T_{\sqrt{1-\epsilon}} f\|_2 \leq \|f\|_p, \quad \textit{where } p = p(q,\epsilon) = 1 + (1-\epsilon)^{(2-4/q)/\log(q-1)}.$$

The second 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 3.5. *Suppose $f : [q]^n \to [0,1]$ has $\mathrm{Inf}_i^{(1-\eta)}[f] \leq \tau \leq (\log q)^{-(\log q)/c}$ for all $i \in [n]$, where $\eta < c(\log q)/\log(1/\tau)$ and $c > 0$ is a certain universal constant. Let $\mu = \mathbf{E}[f]$. Then for any $0 < \epsilon < 1$,*

$$\mathbb{S}_{1-\epsilon}[f] \leq \Gamma_{1-\epsilon}(\mu) + \frac{\log q}{c\epsilon} \cdot \frac{\log\log(1/\tau)}{\log(1/\tau)}.$$

In the above theorem, the quantity $\Gamma_{1-\epsilon}(\mu)$ 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 $\mathbf{Pr}[\boldsymbol{x} \leq t] = \mu$. We will use the following estimate:

Proposition 3.6. *Assume $0 < \epsilon < .1$ and $0 \leq \mu \leq \exp(-1/\sqrt{\epsilon})/\sqrt{\epsilon}$. Then $\Gamma_{1-\epsilon}(\mu) \leq \mu^{1+\epsilon/(2-\epsilon)}$.*

This estimate follows from Corollary 10.2 in [KKMO07]. (The expression in that corollary is in fact an upper bound on $\Gamma_{1-\epsilon}(\mu)$ for all $0 < \epsilon < 1$ and $0 \leq \mu \leq 1/2$, as can be verified using the inequality in Proposition 6.1 of [KKMO07]. The simplified bound $\mu^{1+\epsilon/(2-\epsilon)}$ holds when $\epsilon < .1$ and $\mu \leq \exp(-1/\sqrt{\epsilon})/\sqrt{\epsilon}$.)

3.4 Dictator vs. Small Noisy-Influence Tests

The work of Khot, Kindler, Mossel, and O'Donnell [KKMO07] introduced a now-standard methodology for proving hardness results based on the UGC: namely, the construction of "Dictator vs. Small Noisy-Influences[1] Tests". (We will sometimes call these just Dictator Tests, for brevity.) Formally speaking, a *test* for functions f with domain $[q]^n$ is nothing more than an explicit instance $\mathcal{T}$ of a weighted constraint satisfaction problem with variable set $[q]^n$. Usually, however, it is thought of as a "probabilistic spot-check" on an assignment function f, where one chooses a constraint from $\mathcal{T}$ with probability equal to the constraint's weight, and then "tests" whether f's assignment satisfies the constraint. An important aspect of such a test is the "type of constraint" it uses. Naturally, in this paper we will be considering two-variable linear equation constraints; specifically, testing functions $f : [q]^n \to \mathbb{Z}_m$ using constraints of the form $f(x) - f(y) = c$, where $c \in \mathbb{Z}$.

Before defining Dictator Tests we need to introduce another small technical detail, that of testing *averages* of functions. Given a test for functions $f : [q]^n \to \mathbb{Z}_m$, say, we can think of it more generally as a test for functions $f : [q]^n \to \Delta_m$. To understand this, one should think of a function with range Δ_m as a "randomized" function into $\mathbb{Z}_m$. I.e., to apply the test $\mathcal{T}$ to a function $f : [q]^L \to \Delta_m$, one first chooses a random constraint as usual in $\mathcal{T}$; say it is $f(x) - f(y) = c$. One then chooses $\boldsymbol{a} \sim f(x)$ and $\boldsymbol{b} \sim f(y)$ (independently) and finally, one checks the constraint $\boldsymbol{a} - \boldsymbol{b} = c$.

We may now informally state what a *Dictator vs. Small Noisy-Influences Test* is. It is a test for functions $f : [q]^n \to \Delta_m$ with the following two properties: (i) *Dictator functions* — i.e.,

[1] Actually, they used "low-degree influences", but the distinction is inessential.

functions of the form $f(x) = x_i$ — pass the test with high probability. (Here we are interpreting the integer $x_i \in [q]$ also as an element of $\mathbb{Z}_m$, and thus also as an element of Δ_m.) In other words, $\mathrm{Val}_{\mathcal{T}}(f)$ is large when f is a dictator. (ii) Functions f satisfying $\mathrm{Inf}_i^{(1-\eta)}[f] \leq \tau$ for all $i \in [n]$ pass the test with low probability, where here η and τ should be thought of as very small constants. More formally:

Definition 3.7. Let $\mathcal{T}$ be a test for functions $f : [q]^n \to \Delta_m$. We say that $\mathcal{T}$ has *completeness* at least c if every dictator function $f(x) = x_i$ passes the test with probability at least c. We say that $\mathcal{T}$ has (τ, η)-*soundness* at most s if every function $f : [q]^n \to \Delta_m$ satisfying $\mathrm{Inf}_i^{(1-\eta)}[f] \leq \tau$ for all $i \in [n]$ passes the test with probability at most s. Finally, given a family of tests $(\mathcal{T}_n)$, where $\mathcal{T}_n$ test functions $f : [q]^n \to \Delta_m$, we say it has *soundness* s if for every $\kappa > 0$ there exists $\tau, \eta > 0$ such that each $\mathcal{T}_n$ has (τ, η)-soundness at most $s + \kappa$.

Khot, Kindler, Mossel, and O'Donnell (implicitly) showed the following connection between Unique-Games-hardness and Dictator Tests (see also, e.g., [Rag09, Theorem 7.6]):

Theorem 3.8. *Suppose there exists a completeness-c, soundness-s family $(\mathcal{T}_L)_{L \in \mathbb{N}}$ of Dictator Tests for functions $f : \Omega^L \to \Delta_{|\Sigma|}$, using constraints of type Φ over alphabet Σ. Then assuming the UGC, for all $\epsilon > 0$ it is* NP*-hard to distinguish* Max-Φ *CSP instances $\mathcal{I}$ with $\mathrm{Opt}(\mathcal{I}) \geq c$ from instances with $\mathrm{Opt}(\mathcal{I}) \leq s + \epsilon$.*

Thus the task of proving NP-hardness results for constraint satisfaction problems assuming the UGC is reduced to constructing Dictator Tests. Our main contribution is to construct and analyze an appropriate such test using Max-2Lin-type constraints. Because we are working in a slightly nonstandard setting (testing functions $f : [q]^L \to \mathbb{Z}_m$, where m may be "super-constantly" large as a function of L), we will give the complete reduction based on our Dictator Test in Appendix A.

4 The new Dictator Test

We begin by stating our new family of Dictator vs. Small Noisy-Influences Tests. Given parameters $0 < \epsilon < 1$ and $q \in \mathbb{N}$, we define the following test $\mathcal{T}_{q,\epsilon}$ for functions f with domain $[q]^L$:

Test $\mathcal{T}_{q,\epsilon}$:

- Choose $\boldsymbol{x}, \boldsymbol{x}' \sim [q]^L$ to be a pair of $(1-\epsilon)$-correlated random strings.
- Choose $\boldsymbol{c}, \boldsymbol{c}' \sim [q/\log(q)]$ independently and uniformly.
- Define $\boldsymbol{y} = \boldsymbol{x} \oplus_q (\boldsymbol{c}, \boldsymbol{c}, \ldots, \boldsymbol{c})$, and define $\boldsymbol{y}' = \boldsymbol{x}' \oplus_q (\boldsymbol{c}, \boldsymbol{c}, \ldots, \boldsymbol{c})$.
- Test the constraint " $f(\boldsymbol{y}) - \boldsymbol{c} = f(\boldsymbol{y}') - \boldsymbol{c}'$ " (equivalently, " $f(\boldsymbol{y}) - f(\boldsymbol{y}') = \boldsymbol{c} - \boldsymbol{c}'$ ").

As discussed, one can also think of this test as an explicit weighted constraint satisfaction problem of Max-2Lintype over the variable set $[q]^L$. The constraint $f(\boldsymbol{y}) - \boldsymbol{c} = f(\boldsymbol{y}') - \boldsymbol{c}'$ should be thought of as a formal expression, since we have not yet specified the *range* of the assignment f. In fact, we will analyze the test's properties when the range of f varies over different $\mathbb{Z}_m$'s.

In light of the connection between Unique-Games-hardness and Dictator Tests (Theorem 3.8), to prove our main Theorem 1.9 it will suffice (as we verify in Appendix A) to show the following.

Theorem 4.1. *The Dictator Test $\mathcal{T}_{q,\epsilon}$ uses integer constants c_{ij} in $[-q/\log(q), q/\log(q)]$ and has the following two properties:*

Completeness: *For each $m : q \leq m \leq +\infty$, the L dictator functions $f : [q]^L \to \mathbb{Z}_m$ pass the test $\mathcal{T}_{q,\epsilon}$ with probability at least $1 - \epsilon - O(1/\log(q))$.*

Soundness: *Assume* $0 < \epsilon < .1$ *and that* $q \geq \exp(1/\sqrt{\epsilon})$ *an integer. Assume* $f : [q]^L \to \Delta_m$ *satisfies* $\mathrm{Inf}_i^{(1-\eta)}[f] \leq \tau \leq (\log q)^{-(\log q)/c}$ *for all* $i \in [L]$, *where* $\eta < c(\log q)/\log(1/\tau)$ *(and* c *is the constant from Theorem 3.5). Assume further that* $q/\log(q) \leq m \leq \infty$. *Then* f *passes the test* $\mathcal{T}_{q,\epsilon}$ *with probability less than*

$$\widetilde{O}(1/q)^{\epsilon/(2-\epsilon)} + \frac{\widetilde{O}(q^{\log q})}{\epsilon} \cdot \frac{\log\log(1/\tau)}{\log(1/\tau)}.$$

The Completeness part of Theorem 4.1 is easy to verify:

Proof. Suppose $f(x) = x_j$ for some $j \in [L]$. In the test $\mathcal{T}_{q,\epsilon}$ we have $\boldsymbol{x}_j = \boldsymbol{x}'_j$ except with probability at most ϵ. When the event happens, write $\boldsymbol{b}$ for the common value. We further have that $\boldsymbol{b}$ is at most $q - \lfloor q/\log(q) \rfloor$ except with probability at most $O(1/\log(q))$. Thus with probability at least $1-\epsilon-O(1/\log(q))$ we have both $\boldsymbol{y}_j = \boldsymbol{b}+\boldsymbol{c}$ and $\boldsymbol{y}'_j = \boldsymbol{b}+\boldsymbol{c}'$ *as integers in* $[q]$; i.e., the $\oplus_q$ does not cause "wrap-around". Thus $f(\boldsymbol{y})$ will equal the integer $\boldsymbol{b}+\boldsymbol{c}$ within $\mathbb{Z}_m$, and similarly $f(\boldsymbol{y}')$ will equal $\boldsymbol{b}+\boldsymbol{c}'$ within $\mathbb{Z}_m$, and the tested constraint will be satisfied. □

The next two subsections of the paper are devoted to the proof of the Soundness part of Theorem 4.1. In the first subsection we prove a technical lemma bounding the noise stability of functions $f : [q]^L \to \Delta_m$ which have $\|f_j\|_\infty$ small for each $j \in \mathbb{Z}_m$. In the subsequent subsection, we complete the proof of the soundness of our test.

4.1 Technical lemma

Our soundness analysis relies on the following technical lemma; the crucial aspect of it is that the upper bound we give on the noise stability does not depend on m.

Lemma 4.2. *Fix* $0 < \epsilon < .1$ *and let* $q \geq \exp(1/\sqrt{\epsilon})$ *be an integer. Further, let* $L, m \in \mathbb{N}$ *and* $0 < \eta < 1$. *Assume* $g : [q]^L \to \Delta_m$ *satisfies* $\mathrm{Inf}_i^{(1-\eta)}[g] \leq \tau \leq (\log q)^{-(\log q)/c}$ *for all* $i \in [L]$, *where* $\eta < c(\log q)/\log(1/\tau)$ *(and* c *is the constant from Theorem 3.5).*

Then if $\mathbf{E}_{\boldsymbol{x}}[g(\boldsymbol{x})_a] \leq \log(q)/q$ *for all* $a \in [m]$, *it follows that*

$$\mathbb{S}_{1-\epsilon}[g] < \widetilde{O}(1/q)^{\epsilon/(2-\epsilon)} + \frac{\widetilde{O}(q^{\log q})}{\epsilon} \cdot \frac{\log\log(1/\tau)}{\log(1/\tau)}.$$

Proof. Write $\mu_a = \mathbf{E}_{\boldsymbol{x}}[g(\boldsymbol{x})_a]$. We use two different bounds for $\mathbb{S}_{1-\epsilon}[g_a]$ depending on the magnitude of μ_a. The first bound uses the small noisy-influences of g_a (which are certainly smaller than those of g) and the Majority Is Stablest Theorem (Theorem 3.5), yielding

$$\mathbb{S}_{1-\epsilon}[g_a] \leq \Gamma_{1-\epsilon}(\mu_a) + e(\tau), \qquad e(\tau) := \frac{\log q}{c\epsilon} \cdot \frac{\log\log(1/\tau)}{\log(1/\tau)}.$$

We may also use Proposition 3.6 because $\epsilon < .1$ and $\mu_a \leq \log(q)/q \leq \exp(-1/\sqrt{\epsilon})/\sqrt{\epsilon}$; thus

$$\mathbb{S}_{1-\epsilon}[g_a] \leq \mu_a^{1+\epsilon/(2-\epsilon)} + e(\tau). \tag{2}$$

Our second bound is more useful when μ_a is extremely small; it only needs the hypercontractivity theorem (Theorem 3.4), and not the small noisy-influences condition. The theorem gives

$$\mathbb{S}_{1-\epsilon}[g_a] = \|T_{\sqrt{1-\epsilon}} g_a\|_2^2 \leq \|g_a\|_p^2 = \mathbf{E}[g_a^p]^{2/p} \leq \mathbf{E}[g_a]^{2/p} = \mu_a^{2/p},$$

where $p = 1+(1-\epsilon)^{(2-4/q)/\log(q-1)}$ as in Theorem 3.4. One can check that $2/p \geq 1+\epsilon/(1.9\log q)$ for all $0 < \epsilon < 1$ and $q \geq 3$; hence:

$$\mathbb{S}_{1-\epsilon}[g_a] \leq \mu_a^{1+\epsilon/(1.9\log q)}. \tag{3}$$

We now put the two bounds together:

$$\begin{aligned}
\mathbb{S}_{1-\epsilon}[g] &= \sum_{a \in [m]} \mathbb{S}_{1-\epsilon}[g_a] \\
&= \sum_{a:\mu_a \geq q^{-\log q}} \mathbb{S}_{1-\epsilon}[g_a] + \sum_{a:\mu_a < q^{-\log q}} \mathbb{S}_{1-\epsilon}[g_a] \\
&= \sum_{a:\mu_a \geq q^{-\log q}} (\mu_a^{1+\epsilon/(2-\epsilon)} + e(\tau)) \quad + \sum_{a:\mu_a < q^{-\log q}} \mu_a^{1+\epsilon/(1.9\log q)} \qquad \text{(using (2), (3))}.
\end{aligned}$$

Since g's range is Δ_m we have $\sum_{a\in[m]} \mu_a = 1$. Thus the first sum above is at most

$$q^{\log q} e(\tau) + \sum_{a:\mu_a \geq q^{-\log q}} \mu_a^{1+\epsilon/(2-\epsilon)} \leq q^{\log q} e(\tau) + \max_a \mu_a^{\epsilon/(2-\epsilon)} \leq q^{\log q} e(\tau) + (\log(q)/q)^{\epsilon/(2-\epsilon)}$$

using the assumed upper bound on μ_a. The second sum above is at most

$$\max_{a:\mu_a < q^{-\log q}} \mu_a^{\epsilon/(1.9\log q)} \leq (q^{-\log q})^{\epsilon/(1.9 \log q)} = q^{-\epsilon/1.9}.$$

Thus we conclude

$$\mathbb{S}_{1-\epsilon}[g] \leq q^{\log q} e(\tau) + (\log(q)/q)^{\epsilon/(2-\epsilon)} + q^{-\epsilon/1.9} < \widetilde{O}(1/q)^{\epsilon/(2-\epsilon)} + \frac{\widetilde{O}(q^{\log q})}{\epsilon} \cdot \frac{\log\log(1/\tau)}{\log(1/\tau)}$$

as claimed. □

4.2 Soundness of the test

This section is devoted to the proof of the Soundness part of Theorem 4.1.

Proof. Given f as in the statement of the theorem, we introduce another randomized function $g : [q]^L \to \Delta_m$. Specifically, $g(x)$ is defined to be the distribution function on $\boldsymbol{a} \in \mathbb{Z}_m$ given by the following experiment:

- Choose $\boldsymbol{c} \sim [q/\log(q)]$ uniformly at random.
- Choose $\boldsymbol{b}$ according to the distribution $f(x \oplus_q (\boldsymbol{c}, \boldsymbol{c}, \ldots, \boldsymbol{c}))$.
- Define $\boldsymbol{a} = \boldsymbol{b} - \boldsymbol{c} \in \mathbb{Z}_m$.

Thus in the test $\mathcal{T}_{q,\epsilon}$, once $\boldsymbol{x}$ and $\boldsymbol{x'}$ are chosen the probability that f passes the test is equal to the probability that independent draws from $g(\boldsymbol{x})$ and $g(\boldsymbol{x'})$ yield the same value in $\mathbb{Z}_m$. I.e.,

$$\mathbf{Pr}[f \text{ passes the constraint}] = \mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{x'}}[\langle g(\boldsymbol{x}), g(\boldsymbol{x'})\rangle] = \mathbb{S}_{1-\epsilon}[g].$$

It thus suffices to bound $\mathbb{S}_{1-\epsilon}[g]$.

Our first task is to show that g has small noisy-influences. Define the operator S_c for $c \in \mathbb{Z}_q$ as follows: $S_c h(x) = h(x \oplus_q (c, c, \ldots, c))$. Define the operator R_c for $c \in \mathbb{Z}_m$ as follows: $(R_c h(x))_a = h(x)_{a+c}$, where the sum $a + c$ is within $\mathbb{Z}_m$. Hence by definition,

$$g = \mathop{\text{avg}}_{c\in[q/\log(q)]} \{R_c S_c f\}. \tag{4}$$

In particular, for each $i \in [L]$ we have

$$\mathrm{Inf}_i^{(1-\eta)}[g] \leq \mathop{\text{avg}}_{c\in[q/\log(q)]} \{\mathrm{Inf}_i^{(1-\eta)}[R_c S_c f]\}$$

by the convexity of noisy-influences (Proposition 3.3). But it's easy to see that $\mathrm{Inf}_i^{(1-\eta)}[R_c h] = \mathrm{Inf}_i^{(1-\eta)}[h]$ and $\mathrm{Inf}_i^{(1-\eta)}[S_c h] = \mathrm{Inf}_i^{(1-\eta)}[h]$. Hence we conclude $\mathrm{Inf}_i^{(1-\eta)}[g] \leq \mathrm{Inf}_i^{(1-\eta)}[f] \leq \tau$ for all $i \in [L]$.

We now make the key observation. For $a \in \mathbb{Z}_m$, define $\mu_a = \mathbf{E}_{\boldsymbol{x}\sim[q]^L}[g(\boldsymbol{x})_a]$. Using the original definition of g we have

$$\mu_a = \mathop{\mathbf{Pr}}_{\substack{\boldsymbol{x},\boldsymbol{c}\sim[q/\log(q)],\\ \boldsymbol{b}\sim f(\boldsymbol{x}\oplus_q(\boldsymbol{c},\boldsymbol{c},\ldots,\boldsymbol{c}))}}[\boldsymbol{b}-\boldsymbol{c}=a] = \mathop{\mathbf{E}}_{\boldsymbol{x},\boldsymbol{c}\sim[q/\log(q)]}[f(\boldsymbol{x}\oplus_q(\boldsymbol{c},\boldsymbol{c},\ldots,\boldsymbol{c}))_{a+\boldsymbol{c}}],$$

where the expressions $\boldsymbol{b}-\boldsymbol{c}=a$ and $a+\boldsymbol{c}$ are treated within $\mathbb{Z}_m$. But the joint distribution of $\boldsymbol{c}$ and $\boldsymbol{x}\oplus_q(\boldsymbol{c},\boldsymbol{c},\ldots,\boldsymbol{c})$ is identical to the joint distribution of $\boldsymbol{c}$ and $\boldsymbol{y}$, where $\boldsymbol{y}\sim[q]^L$ is uniform and *independent* of $\boldsymbol{c}$. Hence

$$\mu_a = \mathop{\mathbf{E}}_{\boldsymbol{y},\boldsymbol{c}\sim[q/\log(q)]}[f(\boldsymbol{y})_{a+\boldsymbol{c}}] \le \max_{y\in[q]^L}\left\{\mathop{\mathbf{E}}_{\boldsymbol{c}\sim[q/\log(q)]}[f(y)_{a+\boldsymbol{c}}]\right\} \le \log(q)/q \quad \text{for all } a\in\mathbb{Z}_m, \tag{5}$$

since $\sum_b f(x)_b = 1$ and $m \ge q/\log(q)$.

Having established (5) and also $\mathrm{Inf}_i^{(1-\eta)}[g] \le \tau$ for all i, we may bound $\mathbb{S}_{1-\epsilon}[g]$ and thus complete the proof using the technical Lemma 4.2. (In the case that $m=\infty$ we may still apply the lemma because g's outputs are nonzero on only finitely many coordinates; hence we may consider g's range to be a finite-dimensional simplex.) □

References

[DMR06] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing*, pages 344–353, 2006. 1

[DSC96] Persi Diaconis and Laurent Saloff-Coste. Logarithmic Sobolev inequalities for finite Markov chains. *Annals of Applied Probability*, 6(3):695–750, 1996. 6

[GMR08] Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. Beating the random ordering is hard: Inapproximability of Maximum Acyclic Subgraph. In *Proc. 49th IEEE Symposium on Foundations of Computer Science*, 2008. To appear. 1, 3

[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. 2, 4

[Hås97] Johan Håstad. Some optimal inapproximability results. In *Proc. 29th ACM Symposium on Theory of Computing*, pages 1–10, 1997. 1

[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. 1

[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. 1, 5, 6, 11

[KM10] Subhash Khot and Dana Moshkovitz. Hardness of approximately solving linear equations over reals. *Electronic Colloquium on Computational Complexity*, 5(53), 2010. 2

[KR08] S. Khot and O. Regev. Vertex cover might be hard to approximate to within $2-\epsilon$. *Journal of Computer & System Sciences*, 74(3):335–349, 2008. 1, 11

[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. 6

[Rag09] Prasad Raghavendra. *Approximating NP-hard problems: efficient algorithms and their limits.* PhD thesis, University of Washington, 2009. 1, 5, 6, 7, 12

A The reduction from $\mathsf{Unique\text{-}Games}_L$

In this section we show how to use our Dictator Test to obtain our main UG-hardness result, Theorem 1.9. We reiterate that we are essentially using Theorem 3.8, implicitly proved in [KKMO07]; we give the full deduction here for completeness and because we are working in a slightly nonstandard setting.

For technical convenience, we will use the following equivalent version of the Unique Games Conjecture due to Khot and Regev [KR08, Lemma 3.6]:

Theorem A.1. *Assume the UGC. For all small $\zeta, \gamma > 0$, there exists $L \in \mathbb{N}$ such given an unweighted $\mathsf{Unique\text{-}Games}_L$ instance $\mathcal{G} = (U, V, E, (\pi_{u,v})_{(u,v)\in E})$ which is U-regular, it is NP-hard to distinguish the following two cases:*

1. *There is an assignment $A : (U \cup V) \to [L]$ and a subset $U' \subseteq U$ with $|U'|/|U| \geq 1 - \zeta$ such that A satisfies all constraints incident on U'.*
2. *There is no γ-good assignment A.*

Our main task, which we will carry out in the next subsection, will be to prove the following slight variant of Theorem 1.9, wherein we write $s(q, \epsilon) = \widetilde{O}(1/q)^{\epsilon/(2-\epsilon)}$ for the main term in the Soundness part of Theorem 4.1:

Theorem A.2. *Fix $0 < \epsilon < .1$ rational and $q \geq \exp(1/\sqrt{\epsilon})$ an integer. For any $L \in \mathbb{N}$, there is a polynomial-time reduction mapping non-bipartite, unweighted $\mathsf{Unique\text{-}Games}_L$ instances $\mathcal{G}$ into $\mathsf{Max\text{-}2Lin}$ instances $\mathcal{I}$ having the following properties:*

- *(Completeness.) If statement 1 in Theorem A.1 holds for $\mathcal{G}$, then there is an integer assignment to the variables in $\mathcal{I}$ satisfying at least $(1-\zeta)(1-\epsilon - O(1/\log(q)))$-weight of the equations.*
- *(Soundness.) If there is no γ-good assignment for $\mathcal{G}$ where $\gamma = \gamma(q, \epsilon) > 0$ is sufficiently small, then there is no integer assignment to the variables in $\mathcal{I}$ which satisfies at least $(3s(q,\epsilon))$-weight of the equations modulo m, for any integer $m \geq q/\log(q)$.*

By combining Theorem A.2 with Theorem A.1, taking $\zeta = 1/\log(q)$ and $\gamma = \gamma(q, \epsilon) > 0$ as necessary, we obtain the following variant of Theorem 1.9:

Theorem A.3. *Assume the UGC. For any $0 < \epsilon < .1$ rational and $q \geq \exp(1/\sqrt{\epsilon})$ an integer, the following holds: Given an instance $\mathcal{I}$ of $\mathsf{Max\text{-}2Lin}$ in which the integer constants c_{ij} are in the range $[-q/\log(q), q/\log(q)]$, it is NP-hard to distinguish the following two cases:*

- *There is a $(1 - \epsilon - O(1/\log(q)))$-good integer assignment to the variables.*
- *There is no assignment to the variables which is $\widetilde{O}(1/q)^{\epsilon/(2-\epsilon)}$-good* modulo any integer *$m \geq q/\log(q)$.*

From this, we can deduce our main Theorem 1.9 for ϵ' and δ' by taking ϵ in Theorem A.3 a rational of the form $\epsilon' - \Theta(1/\log(q))$.

A.1 Proof of Theorem A.2

We now prove Theorem A.2.

Proof. The reduction is essentially as in [KKMO07]. Given the $\mathsf{Unique\text{-}Games}_L$ instance $\mathcal{G} = (U, V, E, (\pi_{uv}))$, the reduction produces a weighted $\mathsf{Max\text{-}2Lin}$ instance $\mathcal{I}$ with variable set $V \times [q]^L$. We think of an assignment F to these variables as a collection of functions $f_v : [q]^L \to \mathbb{Z}_m$, one for each $v \in V$. Here we will allow $q/\log(q) \leq m \leq \infty$. For each $u \in V$ we also introduce the randomized function $f_u : [q]^L \to \Delta_m$ defined by

$$f_u(x) = \mathop{\mathbf{E}}_{\boldsymbol{v}:(u,\boldsymbol{v})\in E}[f_{\boldsymbol{v}}^{\pi_{u\boldsymbol{v}}}(x)],$$

where define the functions $f_v^\pi : [q]^L \to \mathbb{Z}_m$ by

$$f_v^\pi(x) = f_v(x \circ \pi^{-1}), \qquad \text{with } x \circ \pi^{-1} \in [q]^L \text{ defined by } (x \circ \pi^{-1})_j = x_{\pi^{-1}(j)}.$$

We now define the instance according to the following probabilistic test:

- Choose $\boldsymbol{u} \in U$ randomly.
- Apply test $\mathcal{T}_{q,\epsilon}$ from Section 4 to $f_{\boldsymbol{u}}$.

Note that by the definition of applying a test to a randomized function, this indeed makes $\mathcal{I}$ a weighted Max-2Lin instance over the variables $V \times [q]^L$. Further, it is easy to check that the reduction from $\mathcal{G}$ to $\mathcal{I}$ thus defined can be carried out in polynomial time assuming ϵ, q, and L are constant.

To prove the Completeness part of Theorem A.2, suppose that assignment A and subset $U' \subseteq U$ are as in statement 1 of Theorem A.1. Define an integer-valued assignment F for $\mathcal{I}$ by taking $f_v(x) = x_{A(v)}$. Then by definition and by the property of A, we will have that $f_u : [q]^L \to \Delta_\mathbb{Z}$ is in fact the $A(u)$th dictator function for all $u \in U'$. Thus by the completeness part of Theorem 4.1, assignment F will pass the test $\mathcal{I}$ with probability at least $\mathbf{Pr}[\boldsymbol{u} \in U'] \cdot (1 - \epsilon - O(1/\log(q))) \geq (1-\zeta)(1 - \epsilon - O(1/\log(q)))$. This finishes the Completeness part of Theorem A.2.

As for the Soundness part of Theorem A.2, choose $\tau = \tau(q,\epsilon) > 0$ small enough so that the error term in the Soundness part of Theorem 4.1 is at most the main term, $s(q,\epsilon)$; choose also $\eta = \eta(q,\epsilon) > 0$ sufficiently small so that the hypothesis therein holds. By way of proving the contrapositive, suppose that there is an integer $m \geq q/\log(q)$ and a $\mathbb{Z}_m$-valued assignment F to $\mathcal{I}$ which passes the test $\mathcal{I}$ with probability at least $3s(q,\epsilon)$. Then by an averaging argument, there must be some subset $U' \subseteq V$ of fractional size at least $s(q,\epsilon)$ such that when $u \in U'$, the test $\mathcal{T}_{q,\epsilon}$ passes f_u with probability at least $2s(q,\epsilon)$. It follows from the Soundness part of Theorem 4.1, along with our choice of τ and η, that

$$\text{for all } u \in U', \quad \exists i_u \in [L] \text{ s.t. } \mathrm{Inf}_{i_u}^{(1-\eta)}[f_u] > \tau. \tag{6}$$

By definition of f_u and by the convexity of noisy-influences (Proposition 3.3) we deduce that for each such $u \in U'$ and $i_u \in [L]$,

$$\tau < \underset{v:(u,v)\in E}{\mathrm{avg}} \left\{ \mathrm{Inf}_{i_u}^{(1-\eta)}[f_v^{\pi_{uv}}] \right\} = \underset{v:(u,v)\in E}{\mathrm{avg}} \left\{ \mathrm{Inf}_{\pi_{uv}(i_u)}^{(1-\eta)}[f_v] \right\}$$
$$\Rightarrow \quad \tau/2 \leq \mathrm{Inf}_{\pi_{uv}(i_u)}^{(1-\eta)}[f_v] \qquad \text{for at least a } \tau/2\text{-fraction of } u\text{'s neighbors } v. \tag{7}$$

For each $v \in V$ let us define

$$C(v) = \{ j \in [L] : \mathrm{Inf}_j^{(1-\eta)}[f_v] > \tau/2 \};$$

thus by (7) we have:

$$\forall\, u \in U', \quad \pi_{uv}(i_u) \in C(v) \text{ for at least a } \tau/2\text{-fraction of } u\text{'s neighbors } v \in V. \tag{8}$$

We claim and will show shortly that $|C(v)| \leq 1/(\eta\tau)$ for all v. Having established this, consider choosing a random assignment $A : (U \cup V) \to [L]$ as follows: for $u \in U'$ set $A(u) = i_u$; for $v \in V$, choose $A(v)$ randomly from $C(v)$ (assuming the set is nonempty); finally, set $A(w)$ arbitrarily in $[L]$ for all unassigned vertices w. Now by (8), for each $u \in U'$ the expected fraction of constraints incident on u which A satisfies is at least $(\tau/2)/(\eta\tau) = \eta\tau^2/2$. Since $|U'|/|U| \geq s(q,\epsilon)$ and $\mathcal{G}$ is U-regular, we conclude that the expected fraction of all constraints in $\mathcal{G}$ that A satisfies is at least $s(q,\epsilon)\eta\tau^2/2$. Taking $\gamma = \gamma(q,\epsilon) = s(q,\epsilon)\eta\tau^2/2$, we conclude that there must exist a γ-good assignment for $\mathcal{G}$.

It remains to verify the claim that $|C(v)| \leq 1/(\eta\tau)$ for all v. We use the following well-known fact (see, e.g., [Rag09, Proposition 3.0.14]):

Fact A.4. *Let* $h : [q]^L \to \mathbb{R}^m$. *Then* $\sum_{j=1}^{L} \mathrm{Inf}_j^{(1-\eta)}[h] \le \mathbf{Var}[h]/(2e\eta)$.

We also need the following small observation: even for arbitrarily large m, if $h : [q]^L \to \Delta_m$ then $\mathbf{Var}[h] \le 1$. This is because $\mathbf{Var}[h] \le \mathbf{E}[\|h\|^2] \le \max_x\{\|h(x)\|^2\} \le 1$, as every point in Δ_m has Euclidean norm at most 1. Thus

$$|C(v)| = \left|\{j \in [L] : \mathrm{Inf}_j^{(1-\eta)}[f_v] > \tau/2\}\right| \le \frac{1/(2e\eta)}{\tau/2} \le 1/(\eta\tau),$$

as claimed. □