On the Power of Unique 2-Prover 1-Round Games

Subhash Khot *
Princeton University, NJ-08544
khot@cs.princeton.edu

ABSTRACT

A 2-prover game is called unique if the answer of one prover uniquely determines the answer of the second prover and vice versa (we implicitly assume games to be one round games). The value of a 2-prover game is the maximum acceptance probability of the verifier over all the prover strategies. We make the following conjecture regarding the power of unique 2-prover games, which we call the Unique Games Conjecture :

The Unique Games Conjecture : For arbitrarily small constants $\zeta,\ \delta > 0$, there exists a constant $k = k(\zeta, \delta)$ such that it is NP-hard to determine whether a unique 2-prover game with answers from a domain of size k has value at least $1 - \zeta$ or at most δ.

We show that a positive resolution of this conjecture would imply the following hardness results :

1. For any $\frac{1}{2} < t < 1$, for all sufficiently small constants $\epsilon > 0$, it is NP-hard to distinguish between the instances of the problem 2-Linear-Equations mod 2 where either there exists an assignment that satisfies $1-\epsilon$ fraction of equations or no assignment can satisfy more than $1 - \epsilon^t$ fraction of equations. As a corollary of the above result, it is NP-hard to approximate the Min-2CNF-deletion problem within any constant factor.

2. For the constraint satisfaction problem where every constraint is the predicate Not-all-equal(a, b, c), $a, b, c \in GF(3)$, it is NP-hard to distinguish between the instances where either there exists an assignment that satisfies $1 - \epsilon$ fraction of the constraints or no assignment satisfies more than $\frac{8}{9} + \epsilon$ fraction of the constraints for an arbitrarily small constant $\epsilon > 0$. We also get a hardness result for a slight variation of approximate coloring of 3-uniform hypergraphs.

We also show that a variation of the Unique Games Conjecture implies that for arbitrarily small constant $\delta > 0$ it is hard to find an independent set of size δn in a graph that is guaranteed to have an independent set of size $\Omega(n)$.

The main idea in all the above results is to use the 2-prover game given by the Unique Games Conjecture as an "outer verifier" and build new probabilistically checkable proof systems (PCPs) on top of it. The uniqueness property plays a crucial role in the analysis of these PCPs.

In light of such interesting consequences, we think it is an important open problem to prove (or disprove) the Unique Games Conjecture. We also present a semi-definite programming based algorithm for finding reasonable prover strategies for a unique 2-prover game. Given a unique 2-prover game with value $1 - \zeta$ and answers from a domain of size k, this algorithm finds prover strategies that make the verifier accept with probability $1 - O(k^2\zeta^{1/5}\sqrt{\log(\frac{1}{\zeta})})$. This result shows that the domain size $k = k(\zeta, \delta)$ must be sufficiently large if the Unique Games Conjecture is true.

*This work was partly supported by Sanjeev Arora's David and Lucile Packard Fellowship and NSF Grant CCR-0098180

STOC'02, May 19-21, 2002, Montreal, Quebec, Canada.

1. INTRODUCTION

The discovery of the PCP Theorem ([4], [3]) and subsequent quantitative improvements in PCP constructions have led to (in many cases optimal) hardness of approximation results for various optimization problems. For example Max-Clique [14], Max-3-SAT [15] and Set Cover [9] to name a few.

However the PCP techniques haven't been successful in obtaining "good" hardness results for some problems like Vertex Cover (see [7] for an exciting new result), Min-2CNF-deletion and coloring of graphs and 3-uniform hypergraphs with a small chromatic number. In this paper we try to identify some promising new directions for attacking these problems.

All PCP constructions today (with the possible exception of [7]) follow the basic paradigm of composing a so called "outer verifier" with an "inner verifier" (proof composition was first introduced by Arora and Safra [4], but the kind of composition we are referring to was first used by Bellare et al [5]). The focus of most of the recent research has been on improving the quality of the inner verifier. Many sophisticated inner verifiers have been constructed (see [14], [15], [21], [13]) based on the Long Codes introduced by Bellare et al [5] and the Fourier Analysis techniques developed by Håstad ([14], [15]). However the outer verifier has remained untouched. All PCP constructions use the same outer verifier, namely the one obtained by parallel repetition of a

2-prover protocol for Gap-3SAT. The soundness property required of the outer verifier is given by the Raz's Parallel Repetition Theorem [20] and we henceforth call this verifier the Raz Verifier.

In this paper, we point out that one promising route for getting good hardness results for problems for which PCP techniques have failed so far, is to construct an outer verifier with "better properties". The Raz Verifier is basically a 2-prover game with the following crucial properties :

1. For arbitrarily small $\delta > 0$, it is NP-hard to determine whether the value of the game is 1 or at most δ.
2. The answers of the provers are from a domain of size k where k is a constant depending on δ.
3. The answer of the second prover uniquely determines the answer of the first prover.

One might expect the property (3) to be even stronger, i.e. the answer of the second prover uniquely determines the answer of the first prover and vice versa. In fact such games have been considered in literature before ([10], [8]) and they are called "unique games". However, to the best of our knowledge, the question whether unique 2-prover games (with $(1-\epsilon,\ \delta)$ gap in their value) are powerful enough to capture NP hasn't been considered before. This question is precisely the focus of this paper and we make the following (rather bold) conjecture :

The Unique Games Conjecture : For arbitrarily small constants $\zeta,\ \delta > 0$, there exists a constant $k = k(\zeta, \delta)$ such that it is NP-hard to determine whether a unique 2-prover game with answers from a domain of size k has value at least $1-\zeta$ or at most δ.

An important point here is that one can trivially determine whether a unique 2-prover game has value 1. Therefore the gap in the above conjecture is $(1-\zeta, \delta)$ as opposed to the gap $(1, \delta)$ in the Raz Verifier. In other words, NP-hard unique games must lose perfect completeness.

We show that a positive resolution of this conjecture would have many interesting consequences. We use the 2-prover game given by the Unique Games Conjecture as an outer PCP verifier and build appropriate inner verifiers to prove the following results :

1. For every $\frac{1}{2} < t < 1$, for all sufficiently small $\epsilon > 0$, it is NP-hard to distinguish between instances of 2-Linear-Equations mod 2, where either there exists an assignment satisfying at least $1-\epsilon$ fraction of equations or no assignment satisfies more than $1-\epsilon^t$ fraction of equations.

   This result is essentially due to Håstad [16]. He proposed a so called "codeword test" for testing Long Codes and analyzed it using Bourgain's theorem [6] on Fourier spectrum of boolean functions. However he wasn't able to give a "consistency test" which would work for the Raz Verifier. The (minor) contribution of this paper is to show that if one uses the outer verifier given by the Unique Games Conjecture, it is indeed possible to construct and analyze a consistency test implying the above hardness result.

   This hardness result is tight since the algorithm of Goemans and Williamson [12] for 2-Linear-Equations mod 2, on an instance with optimum $1-\epsilon$ produces a solution with value $1-O(\sqrt{\epsilon})$.

2. A simple reduction from 2-Linear-Equations to 2-SAT gives a similar result, i.e. a $(1-\epsilon, 1-\epsilon^t)$ gap for 2-SAT for any $\frac{1}{2} < t < 1$. As a corollary, it is NP-hard to approximate Min-2SAT-deletion (also called Min-2CNF-deletion) within any constant factor. On the algorithmic side, Zwick's algorithm [23], on a 2-SAT instance with optimum $1-\epsilon$ produces an assignment with value $1-O(\epsilon^{1/3})$. Klein et al [19] give $O(\log n \log\log n)$ approximation for Min-2CNF-deletion.

3. Guruswami et al [13] (also see [18]) show that the constraint satisfaction problem associated with the predicate Not-all-equal(a,b,c,d) where a,b,c,d are binary variables, is hard to approximate better than a random assignment. They use this fact to derive hardness results for 4-uniform hypergraph coloring. However their techniques do not work for 3-uniform hypergraphs and one of the reasons is that for the predicate Not-all-equal(a,b,c) over binary variables, there does exist an algorithm that does better than a random assignment [22].

   However we show that the Unique Games Conjecture implies that the predicate Not-all-equal(a,b,c) over *ternary* variables is hard to approximate better than a random assignment. We also derive hardness result for a variation of 3-uniform hypergraph coloring which we call "semi-coloring". In this problem we are given a 3-uniform hypergraph and the goal is to color the vertices so that $1-\eta$ fraction of the edges are non-monochromatic (as opposed to *all* edges non-chromatic) where η is a given parameter. (the need for considering this version of coloring is due to the inherent loss of perfect completeness in the Unique Games Conjecture). We show that it is NP-hard to semi-color a 3-semi-colorable 3-uniform hypergraph with constantly many colors.

4. Frieze and Jerrum [11] give an algorithm for Max-k-cut that achieves a factor roughly $1-\frac{1}{k}+\frac{2\ln k}{k^2}$. There is an almost-matching hardness result by Kann et al [17] who show a hardness factor of $1-\frac{1}{34k}$ for this problem. However in their reduction the value of the maximum k-cut in the completeness case is $1-\Omega(\frac{1}{k})$ which is bounded away from 1. It is an interesting open problem whether a similar hardness result holds with perfect completeness or near-perfect completeness.

   We show that for any $t > \frac{1}{2}$, for all sufficiently large constants k, it is NP-hard to distinguish between the instances of Max-k-cut where the optimum value of a k-cut is either $1-\zeta$ or at most $1-\frac{1}{k(\log k)^t}$ where $\zeta > 0$ is an arbitrarily small constant.

We also consider the following relaxation of the uniqueness property. We say that a 2-prover game has "d-to-1 property" if the answer of the second prover uniquely determines the answer of the first prover and for every answer of the first prover, there are at most d answers for the second prover for which the verifier would accept. We assume d to be a fixed integer and $d \geq 2$. Consider the following conjecture :

d-to-1 Conjecture : For arbitrarily small constant $\delta > 0$, there exists a constant $k = k(\delta)$ such that it is NP-hard

to determine whether a 2-prover game with d-to-1 property and answers from a domain of size at most k has value 1 or at most δ.

Note that in contrast with the Unique Games Conjecture, we can hope for perfect completeness in the d-to-1 Conjecture (since $d \geq 2$). We use some of the techniques from Dinur and Safra's paper [7] to show that the d-to-1 Conjecture implies the following results :

1. For arbitrarily small $\epsilon,\ \delta > 0$, it is hard to find an independent set of size δn in a graph which is guaranteed to have an independent set of size $(1 - \frac{1}{2^{1/d}} - \epsilon)n$. (see [1] for an algorithmic result). Note that Dinur and Safra's result [7] does not imply such a result for independent sets. Such a result is equivalent to the existence of a PCP with zero free bits, completeness $\Omega(1)$ and arbitrarily low soundness, which is an open problem.

2. From the above result it follows that if 2-to-1 Conjecture is true, it would imply $\sqrt{2} - \epsilon$ hardness for Vertex Cover which is better than the factor 1.3606 by Dinur and Safra. In fact, Dinur and Safra do use an analog of 2-to-1 property. We do not elaborate on this due to space limitations.

In light of such interesting consequences of the Unique Games Conjecture, we think it is an important open problem to prove or disprove it. In this paper, we also present a semidefinite programming based algorithm giving the following theorem :

THEOREM 1. *There exists a (poly-time) algorithm such that given a unique 2-prover game with value $1 - \epsilon$ and answers from a domain of size k, it finds prover strategies that make the verifier accept with probability $1 - O(k^2 \epsilon^{1/5} \sqrt{\log(\frac{1}{\epsilon})})$.*

Andersson et al [2] proved a similar result for the problem 2-Linear-Equations mod p, where the constraints are linear equations mod p with every equation containing exactly 2 variables. Such constraints have the *uniqueness property* since the value to one variable in the equation uniquely determines the value to the second variable. Our algorithm is simpler and more general than that of Andersson et al.

Theorem 1 shows that if at all the Unique Games Conjecture is true, the domain size required $k = k(\zeta, \delta)$ must be at least $\frac{1}{\zeta^{1/10}}$. A trivial bound $k \geq \frac{1}{\delta}$ also holds, since the provers can choose their answers uniformly at random from the domain of possible answers and satisfy the verifier with probability at least $\frac{1}{k}$.

Overview of the paper : Section 2 provides the preliminary background. We prove the results for 2-Linear-Equations mod 2 and Min-2CNF-deletion in Section 3. We prove the hardness of predicate Not-all-equal(a, b, c) over ternary variables in Section 4. We prove Theorem 1 in Section 5 and appendix A. Proofs of all the other results are omitted from this extended abstract since they are quite lengthy and involved. Section 6 concludes with a few remarks as to why it would be difficult to either prove or disprove the Unique Games Conjecture.

2. PRELIMINARIES

This section gives a preliminary background on PCPs, 2-prover games, Long Codes and the basic paradigm of PCP constructions.

2.1 Probabilistically Checkable Proofs

A language L is said to have a probabilistic checkable proof system with parameters (r, q, c, s) if there exists a probabilistic polynomial time verifier which on input x of size n and a proof Π,

- Uses $r = r(n)$ random bits and queries $q = q(n)$ bits from the proof Π.
- Depending on the bits read from the proof it accepts or rejects.
- It has the following two properties :

  (Completeness) : If $x \in L$, there exists a proof Π which the verifier accepts with probability $\geq c$.

  (Soundness) : If $x \notin L$, the verifier accepts *any* proof with probability at most s.

The parameters $1 \geq c > s > 0$ are called completeness and soundness parameters respectively. We recall the PCP Theorem ([4], [3]) that every language in NP has a PCP system with $c = 1, s = \frac{1}{2}$ and the verifier uses $O(\log n)$ random bits and queries only a constant number of bits from the proof.

2.2 2-Prover 1-Round Games

Consider the following game between 2 provers and a verifier. There is a set V of all possible "questions" that the verifier can ask the first prover and a set of questions W that the verifier can ask the second prover.

A "strategy" of the first prover is a map $L_V : V \to N$ where N is a set of possible answers of the first prover. On a question $v \in V$, the prover returns an answer $L_V(v)$ to the verifier. Similarly the strategy of the second prover is a map $L_W : W \to M$ where M is the set of his possible answers.

The "acceptance predicate" of the verifier is a map

$$\Gamma : V \times N \times W \times M \to \{TRUE, FALSE\}$$

The game works in the following way. The verifier picks a pair of questions (v, w), $v \in V,\ w \in W$ with a certain probability distribution on the set of all pairs. He asks question v to the first prover and the question w to the second prover who return answers $L_V(v)$ and $L_W(w)$ respectively. The verifier accepts iff

$$\Gamma(v, L_V(v), w, L_W(w)) = TRUE$$

The value of the game is defined as the maximum, over all possible prover strategies, of the acceptance probability of the verifier.

We will be interested in games where the answer of the second prover uniquely determines the answer of the first prover, i.e. for every question pair (v, w) asked by the verifier and every answer $b \in M$ of the second prover, there is a unique answer $a \in M$ such that the verifier accepts. In this case, we can associate a function $\pi_{vw} : M \to N$ for every pair (v, w) so that the verifier accepts iff

$$\pi_{vw}(L_W(w)) = L_V(v)$$

A game is called "unique" (see [10], [8]) if $M = N$ and every function π_{vw} is a bijection, i.e. the answer of the second prover uniquely determines the answer of the first prover and vice versa.

Remarks : (1) We assume throughtout this paper that the sets of answers M and N are of constant size. (2) The definition of the unique games differs slightly in ([10], [8]). In their definition, for every answer of one prover, there is *at most one* possible answer of the other prover and vice versa.

2.3 The Label Cover Problem

We define a problem called Label Cover which is equivalent to 2-prover games with the property that the answer of the second prover uniquely determines the answer of the first prover. For the sake of convenience, we prefer to talk in terms of the Label Cover problem instead of 2-prover games.

Definition 1. A Label Cover problem $\mathcal{L}$ consists of a complete bipartite graph $G(V, W)$, with bipartition V, W. An edge (v, w) has a weight p_{vw} with $\sum_{v,w} p_{vw} = 1$. Every vertex in V is supposed to get a label from a set N and every vertex in W is supposed to get a label from a set M. With every edge (v, w) there is associated a "projection" $\pi_{vw} : M \to N$. For an assignment of labels to the vertices of the graph, that is for functions $L_V : V \to N$, $L_W : W \to M$, an edge (v, w) is said to be satisfied if $\pi_{vw}(L_W(w)) = L_V(v)$. The goal is to find an assignment of labels that maximizes the total weight of the satisfied edges. We define $OPT(\mathcal{L})$ to be the maximum weight of edges satisfied by any labeling. A Label Cover problem is called "unique" if $M = N$ and every projection $\pi_{v,w} : M \to M$ is a bijection (i.e. a permutation).

Clearly, a Label Cover problem is same as a 2-prover game where V, W are sets of questions the verifier can ask the two provers and N, M are sets of answers by the provers respectively.

The following theorem is a consequence of the PCP Theorem ([4], [3]) and Raz's Parallel Repetition Theorem [20]. It can be found in any of the papers ([5], [15], [13]).

THEOREM 2. *For every constant $\delta > 0$, there exists a constant $k = k(\delta)$ such that it is NP-hard to determine whether a Label Cover problem $\mathcal{L}$ with answers from sets of size at most k (i.e. $|M|, |N| \leq k$) has $OPT(\mathcal{L}) = 1$ or $OPT(\mathcal{L}) \leq \delta$.*

Remark : It turns out that in the reduction given by Theorem 2, we have $|M| \gg |N|$ and the projections $\pi_{v,w} : M \to N$ are highly many-to-one (this many-to-one-ness increases as δ decreases). The PCP constructions in this paper do not work for such projections. Our constructions need a very stringent condition that the projections be bijections or d-to-1 for some fixed d independent of δ.

It is clear that the Unique Label Cover problem corresponds to a unique 2-prover game. Hence the Unique Games Conjecture can be restated as :

Unique Games Conjecture : For arbitrarily small constants ζ, $\delta > 0$, there exists a constant $k = k(\zeta, \delta)$ such that it is NP-hard to determine whether a unique Label Cover instance with the label sets of size k (i.e. $|M| = k$) has optimum at least $1 - \zeta$ or at most δ.

2.4 Constructing PCPs, Long Codes and Fourier Analysis

We briefly explain a basic paradigm for PCP constructions (see [5], [14], [15], [21], [13]). The verifier can be conceptually divided into an "outer" part and an "inner" part.

The verifier reduces an arbitrary language in NP to a gap-version of Label Cover instance $\mathcal{L}$ as given by Theorem 2. This is called the "outer" part of the verifier.

The verifier then expects the proof to contain "Long Codes" of the labels of vertices in the instance $\mathcal{L}$. The verifier picks some edge(s) of the instance $\mathcal{L}$ and performs some local checks on the supposed long codes of the supposed labels of the endpoints of these edge(s). This local checking is called the "inner" part of the verifier.

For proving the soundness property of the verifier, one shows that if the verifier accepts the encoded proof with "good" probability, then the proof can be "decoded" to define labels for the Label Cover instance $\mathcal{L}$ with a "good" value of $OPT(\mathcal{L})$. This gives a contradiction provided we started with an instance $\mathcal{L}$ with sufficiently small value of $OPT(\mathcal{L})$. Theorem 2 guarantees that $OPT(\mathcal{L})$ can be made arbitrarily small. The proof of the soundness of the verifier relies on the Fourier analysis of the Long Codes.

We define the Long Codes in the following.

Definition 2. A binary Long Code on a set of labels M is indexed by all functions $f : M \to \{-1, 1\}$. The long code A of a label $a \in M$ is given by

$$A(f) = f(a) \qquad \forall\, f : M \to \{-1, 1\}$$

A cheating proof might contain an arbitrary string/table A instead of a correct Long Code. Such tables are handled by their Fourier expansion (see [15] for a detailed exposition)

$$A = \sum_{\alpha \subseteq M} \hat{A}_\alpha \chi_\alpha(f) \qquad \text{where} \quad \chi_\alpha(f) = \prod_{x \in \alpha} f(x)$$

The Fourier coefficients $\widehat{A}_\alpha$ satisfy the Parseval's identity, $\sum_\alpha \widehat{A}_\alpha^2 = 1$.

3. HARDNESS OF 2-LINEAR-EQUATIONS MOD 2

In this section we present a proof of the following theorem.

THEOREM 3. *The Unique Games Conjecture implies that for every $\frac{1}{2} < t < 1$, for all sufficiently small constants $\epsilon > 0$, it is NP-hard to distinguish between the instances of 2-Linear-Equations mod 2, where the fraction of satisfied equations is at least $1 - \epsilon$ or at most $1 - \epsilon^t$.*

This result is essentially due to Håstad [16]. He proposed a test for checking a long code and analyzed it using Bourgain's recent theorem [6] on Fourier spectrum of boolean functions, which itself was inspired by a question raised by Håstad.

The (minor) contribution of this paper is to introduce the Unique Games Conjecture and to show that Håstad's test can be extended to test the consistency between two long codes, giving a PCP verifier that makes a linear test on 2 query bits, has completeness $1 - \epsilon$ and soundness $1 - \epsilon^t$.

Following the standard paradigm, the PCP verifier takes the gap-version of the unique Label Cover problem $\mathcal{L}$ guaranteed by the Unique Games Conjecture and expects the

proof to contain, for every vertex $v \in V$, the long code of the label $L_V(v)$ and for every vertex $w \in W$, the long code of the label $L_W(w)$. These long codes are assumed to be folded, i.e. $A(-f) = -A(f)$ (see [15]).

The verifier picks some edges and checks that the labels along these edges satisfy the corresponding bijections. There is a technical issue of how the edges are picked. Let $p_v = \sum_w p_{vw}$. That is if an edge is picked with a probability equal to its weight, p_v is the probability that the left endpoint is v. Let $\Psi_v : W \to [0,1]$ be defined as $\Psi_v(w) = \frac{p_{vw}}{p_v}$. That is $\Psi_v(w)$ is the conditional probability that the right endpoint of an edge is w given that the left endpoint is v.

Action of the verifier :

1. Pick $v \in V$ with probability p_v. Let A be the (supposed) long code of the (supposed) label of v.
2. Pick a random function $f : M \to \{-1,1\}$ and a "perturbation function" $\mu : M \to \{-1,1\}$. For each $x \in M$, $\mu(x) = 1$ with probability $1-\epsilon$ and $\mu(x) = -1$ with probability ϵ.
3. With probability $\frac{1}{2}$ each, select one of the following actions :
   (a) (Codeword test) Accept iff $A(f) = A(f\mu)$
   (b) (Consistency test) Pick a vertex $w \in W$ with the distribution Ψ_v. Let B be the (supposed) long code of the (supposed) label of w and $\pi = \pi_{vw} : M \to M$ be the bijection between v and w. Accept iff
   $$A(f) = B(f \circ \pi)$$
   where $f \circ \pi$ denotes the composition of functions.

Remark : Håstad proposed and analyzed the codeword test. We propose the consistency test and show that Håstad's analysis can be extended to check consistency provided the Unique Games Conjecture is true.

3.1 Completeness

It is easy to see that the completeness of the test is $1-\frac{\zeta+\epsilon}{2}$ where the outer label cover instance has completeness $1-\zeta$. The test may fail due to 2 reasons : (1) The edge (v,w) picked by the verifier may be an unsatisfied edge of the label cover instance which happens with probability ζ. In this case, the consistency test fails. (2) In a correct proof A is a long code of some $a \in M$. The codeword test fails when $\mu(a) = -1$ which happens with probability ϵ.

The claim about the completeness follows. Note that by the Unique Games Conjecture, ζ can be assumed to be arbitrarily small.

3.2 Soundness Analysis

We use the following (deep) theorem of Bourgain [6].

THEOREM 4. *Let A be any boolean function (for instance a supposed long code) and $k > 0$ an integer. Then for every $\frac{1}{2} < t < 1$, there exists a constant $c_t > 0$ such that*

$$\text{If } \sum_{\alpha \,:\, |\alpha| > k} \widehat{A}_\alpha^2 < c_t k^{-t} \quad \text{then} \quad \sum_{\alpha \,:\, |\widehat{A}_\alpha| \le \frac{1}{10} 4^{-k^2}} \widehat{A}_\alpha^2 < \frac{1}{100}$$

The probability of acceptance of the verifier is clearly

$$\Pr[Acc] = \frac{1}{2} E_{v,f,\mu}\left[\frac{1 + A(f)A(f\mu)}{2} + E_w\left[\frac{1 + A(f)B(f\circ\pi)}{2}\right]\right]$$

Using the Fourier expansion $A = \sum_\alpha \widehat{A}_\alpha \chi_\alpha$ we get

$$E_{f,\mu}[A(f)A(f\mu)] = E_{f,\mu}[\sum_{\alpha_1,\alpha_2} \widehat{A}_{\alpha_1}\widehat{A}_{\alpha_2}\chi_{\alpha_1}(f)\chi_{\alpha_2}(f)\chi_{\alpha_2}(\mu)]$$

Note that α_1, α_2 are subsets of M. We have

$$\chi_{\alpha_1}(f)\chi_{\alpha_2}(f) = \prod_{x\in\alpha_1} f(x) \prod_{x\in\alpha_2} f(x) = \prod_{x \in \alpha_1 \Delta \alpha_2} f(x)$$

where $\alpha_1 \Delta \alpha_2$ is the symmetric difference between the sets α_1 and α_2. The expectation over f is non-zero only if $\alpha_1 \Delta \alpha_2 = \emptyset$, i.e. $\alpha_1 = \alpha_2 = \alpha$. Also $E_\mu[\chi_\alpha(\mu)] = (1-2\epsilon)^{|\alpha|}$. Hence

$$E_{f,\mu}[A(f)A(f\mu)] = \sum_\alpha \widehat{A}_\alpha^2 (1-2\epsilon)^{|\alpha|}$$

Using the Fourier expansion $B = \sum_\beta \widehat{B}_\beta \chi_\beta$, we have

$$E_{f,\mu}[A(f)B(f\circ\pi)] = E_{f,\mu}[\sum_{\alpha,\beta} \widehat{A}_\alpha \widehat{B}_\beta \chi_\alpha(f)\chi_\beta(f\circ\pi)\chi_\beta(\mu)] \tag{1}$$

We have

$$\chi_\beta(f\circ\pi) = \prod_{x\in\beta} f(\pi(x)) = \prod_{y\in\pi(\beta)} f(y) = \chi_{\pi(\beta)}(f)$$

Substituting this in (1) and taking expectation over f we see that the expectation is non-zero only if $\alpha = \pi(\beta)$. Since π is a bijection, $\beta = \pi^{-1}(\alpha)$. Thus (1) can be written as

$$E_{f,\mu}[A(f)B(f\circ\pi)] = \sum_\alpha \widehat{A}_\alpha \widehat{B}_{\pi^{-1}(\alpha)} (1-2\epsilon)^{|\alpha|}$$

Hence the probability of acceptance is

$$\begin{aligned}\Pr[Acc] &= \frac{1}{2} + \frac{1}{4} E_v\left[\sum_\alpha \widehat{A}_\alpha^2 (1-2\epsilon)^{|\alpha|} + \sum_\alpha \widehat{A}_\alpha E_w\left[\widehat{B}_{\pi^{-1}(\alpha)}\right]\right] \\ &= \frac{1}{2} + \frac{1}{4} E_v[R_v + T_v]\end{aligned}$$

If this probability is $\ge 1 - \frac{1}{8} c_t \epsilon^t$ where c_t is as in Theorem 4, we have $E_v[R_v + T_v] \ge 2 - \frac{1}{2} c_t \epsilon^t$. This implies that over the choice of v, with probability at least $\frac{1}{2}$, $R_v + T_v \ge 2 - c_t\epsilon^t$. Fix any such "good" v. We have $R_v \ge 1 - c_t\epsilon^t$ and $T_v \ge 1 - c_t\epsilon^t > \frac{1}{2}$.

$$1 - c_t\epsilon^t \le R_v \le \sum_{\alpha \,:\, |\alpha| \le \epsilon^{-1}} \widehat{A}_\alpha^2 + e^{-2} \sum_{\alpha \,:\, |\alpha| > \epsilon^{-1}} \widehat{A}_\alpha^2$$

$$\Longrightarrow \sum_{\alpha \,:\, |\alpha| > \epsilon^{-1}} \widehat{A}_\alpha^2 < c_t \epsilon^t \tag{2}$$

Taking $k = \epsilon^{-1}$ in Theorem 4, we get

$$\sum_{\alpha \,:\, |\widehat{A}_\alpha| \le \frac{1}{10} 4^{-k^2}} \widehat{A}_\alpha^2 < \frac{1}{100} \tag{3}$$

Now we use the fact that $T_v > \frac{1}{2}$. Call α "good" if $\alpha \subseteq M$ is nonempty, $|\alpha| \le \epsilon^{-1}$ and $|\widehat{A}_\alpha| \ge \frac{1}{10} 4^{-k^2}$. We will show that the contribution of bad α's to T_v is small. First of all,

since the tables are folded, $\widehat{A}_\alpha = 0$ when $|\alpha|$ is even (see [15]). In particular $\widehat{A}_\alpha = 0$ when α is empty. Also

$$\left| \sum_{\alpha \,:\, |\alpha| > \epsilon^{-1}} \widehat{A}_\alpha E_w[\widehat{B}_{\pi^{-1}(\alpha)}] \right| \leq$$

$$\sqrt{\sum_{\alpha \,:\, |\alpha| > \epsilon^{-1}} \widehat{A}_\alpha^2} \sqrt{\sum_\alpha \left| E_w[\widehat{B}_{\pi^{-1}(\alpha)}] \right|^2} \leq$$

$$\sqrt{\sum_{\alpha \,:\, |\alpha| > \epsilon^{-1}} \widehat{A}_\alpha^2} < \sqrt{c_t \epsilon^t}$$

where we used (2). Similarly we use (3), and show that the contribution of α's such that $|\widehat{A}_\alpha| \leq \frac{1}{10} 4^{-k^2}$ to T_v is at most $\frac{1}{10}$. This implies that T_v when restricted to good α's, still remains at least $\frac{1}{4}$. We have

$$E_w \left[\sum_\alpha \widehat{A}_\alpha^2 \widehat{B}^2_{\pi^{-1}(\alpha)} \frac{1}{|\alpha|} \right] \geq \epsilon \; E_w \left[\sum_{\alpha \text{ good}} \widehat{A}_\alpha^2 \widehat{B}^2_{\pi^{-1}(\alpha)} \right] \tag{4}$$

$$\begin{aligned}
&\geq \quad \epsilon \, \tfrac{1}{100} 4^{-2k^2} E_w \left[\textstyle\sum_{\alpha \text{ good}} \widehat{B}^2_{\pi^{-1}(\alpha)} \right] \\
&\geq \quad \epsilon \, \tfrac{1}{100} 4^{-2k^2} E_w \left[\left| \textstyle\sum_{\alpha \text{ good}} \widehat{A}_\alpha \widehat{B}_{\pi^{-1}(\alpha)} \right|^2 \right] \\
&\geq \quad \epsilon \, \tfrac{1}{100} 4^{-2k^2} \left| E_w \left[\textstyle\sum_{\alpha \text{ good}} \widehat{A}_\alpha \widehat{B}_{\pi^{-1}(\alpha)} \right] \right|^2 \\
&\geq \quad \epsilon \, \tfrac{1}{100} 4^{-2k^2} \tfrac{1}{16}
\end{aligned}$$

The expression on the second-last line is just T_v restricted to good α's which we showed to be at least $\frac{1}{4}$. Note that we are assuming that v is good itself, which holds with probability $\frac{1}{2}$.

Now we define a labeling for the Label Cover instance as follows : For a good vertex $v \in V$, pick α with probability $\widehat{A}_\alpha^2$, pick a random element of α and define it to be the label of v. For any vertex $w \in W$, pick β with probability $\widehat{B}_\beta^2$, pick a random element of β and define it to be the label of w.

It is easy to see that the weight of the edges satisfied by this labeling equals the expression (4). Label of v will be defined to be a random element $x \in \alpha$ and the label of w will be defined to be a random element $y \in \pi^{-1}(\alpha)$. With probability $\frac{1}{|\alpha|}$ it holds that $\pi(y) = x$ and the edge (v, w) in the Label Cover instance is satisfied.

Since the expression (4) is at least $\Omega(\epsilon 4^{-2k^2})$, we get a labelling that satisfies edges of total weight $\Omega(\epsilon 4^{-2k^2})$. However this contradicts the fact that $OPT(\mathcal{L}) \leq \delta$ if δ was chosen sufficiently small (see the Unique Games Conjecture). This shows that the soundness is at most $1 - \frac{1}{8} c_t \epsilon^t$ where $t > \frac{1}{2}$ is arbitrary, proving Theorem 3.

Remark : A simple gadget $(x \oplus y = 0 \mapsto \overline{x} \vee y, \; x \vee \overline{y})$ reduces 2-Linear-Equations to 2-SAT and implies a $(1 - \epsilon, 1 - \epsilon^t)$ gap for 2-SAT for any $t > \frac{1}{2}$.

4. HARDNESS OF THE PREDICATE NOT-ALL-EQUAL(A,B,C), A,B,C $\in GF(3)$

In this section we will show hardness of the predicate Not-all-equal(a,b,c) over $GF(3)$. This predicate is TRUE iff a, b, c do not all have the same value. We will prove that

THEOREM 5. *If the Unique Games Conjecture is true, then the following holds : for a constraint satisfaction problem with all constraints of the form Not-all-equal(a, b, c) and the variables from a ternary alphabet, it is NP-hard to determine whether there exists an assignment that satisfies $1 - \epsilon$ fraction of the constraints or no assignment satisfies more than $\frac{8}{9} + \epsilon$ fraction of the constraints, where $\epsilon > 0$ is an arbitrarily small consatnt.*

We will construct a PCP that reads 3 symbols from a proof over ternary alphabet, accepts iff the 3 symbols are not all equal, has completeness $1 - \epsilon$ and soundness $\frac{8}{9} + \epsilon$.

We use Long Code over $GF(3)$ on the set of labels M. Such a code is indexed by all functions $f : M \to \{1, \omega, \omega^2\}$ where ω is the cube root of unity. The Long Code A of $a \in M$ is defined as $A(f) = f(a)$. The Fourier expansion in this setting is

$$A(f) = \sum_\alpha \widehat{A}_\alpha \chi_\alpha(f) \qquad \text{where} \quad \chi_\alpha(f) = \prod_{x \in M} f(x)^{\alpha(x)}$$

and α ranges over all functions $\alpha : M \to GF(3)$. The Fourier coefficient $\widehat{A}_\alpha$ is given by

$$\widehat{A}_\alpha = \frac{1}{3^{|M|}} \sum_{f : M \to \{1, \omega, \omega^2\}} A(f) \overline{\chi_\alpha(f)}$$

Remark : In a correct Long code A, we will have $A(\omega f) = \omega A(f)$ and we may want to force this condition on every (supposed) Long code in the proof. This is called "folding" in the PCP literature. However the specific nature of the predicate Not-all-equal(a, b, c) forbids us from doing so and this makes the analysis more difficult. See [13] for a detailed discussion on this issue.

The verifier is given a unique Label Cover instance $\mathcal{L}$ guaranteed by the Unique Games Conjecture. It expects as a proof the Long codes of the labels of all the vertices in $\mathcal{L}$. The verifier works as follows :

1. Pick a vertex $v \in V$ with probability p_v.
2. Pick 3 vertices w_1, w_2, w_3, each of them independently from the distribution Ψ_v. Let A, B, C be the (supposed) long codes of the (supposed) labels of the vertices w_1, w_2, w_3 respectively. Let $\pi = \pi_{vw_1}$, $\pi' = \pi_{vw_2}$, $\pi'' = \pi_{vw_3}$ be the respective projections.
3. Pick two random functions $f, g : M \to \{1, \omega, \omega^2\}$.
4. Pick a function $\mu : M \to \{\omega, \omega^2\}$ by defining for each $x \in M$, $\mu(x) = \omega$ with probability $\frac{1}{2}$ and $\mu(x) = \omega^2$ with probability $\frac{1}{2}$.
5. Accept iff

$$\text{Not-all-equal}(A(f \circ \pi),\; B(g \circ \pi'),\; C(((\overline{f}\,\overline{g}) \circ \pi'') \cdot \mu))$$

The completeness is $1 - 3\zeta$ where $1 - \zeta$ is the completeness of the outer label cover instance. The verifier picks 3 edges and each of them can be an unsatisfied edge of the Label Cover instance with probability ζ. If all the 3 edges are satisfied, A, B, C are the long codes of some $a, b, c \in M$ respectively and $\pi(a) = \pi'(b) = \pi''(c) = d$ for some $d \in M$. Thus

$A(f \circ \pi) = f(\pi(a)) = f(d)$, $\quad B(g \circ \pi') = g(\pi'(b)) = g(d)$ and

$$C((\overline{f}\overline{g}) \circ \pi'' \cdot \mu) = (\overline{f}\overline{g})(\pi''(c)) \cdot \mu(c) = \overline{f(d)g(d)}\mu(c)$$

and not all three can be equal since $\mu()$ takes values only in the set $\{\omega, \omega^2\}$.

4.1 Soundness Analysis

The following lemma is easily proven.

LEMMA 1. *Let $x, y, z \in \{1, \omega, \omega^2\}$. Then the expression*

$$1 - \frac{1}{9} \sum_{\substack{r_1,r_2,r_3 \in GF(3) \\ r_1+r_2+r_3=0}} x^{r_1} y^{r_2} z^{r_3}$$

equals 0 if $x = y = z$ and 1 otherwise.

From this lemma it is clear that the expression

$$1 - \frac{1}{9} \sum_{\substack{r_1,r_2,r_3 \in GF(3) \\ r_1+r_2+r_3=0}} A(f \circ \pi)^{r_1} B(g \circ \pi')^{r_2} C((\overline{f}\overline{g}) \circ \pi'' \cdot \mu)^{r_3}$$

equals 1 if the test accepts and 0 otherwise. Hence the acceptance probability of the verifier is equal to the expectation of this expression over the choice of $(v, w_1, w_2, w_3, f, g, \mu)$. Let us consider this expectation for a fixed v. We divide the terms in the summation into 3 cases and consider the expectation of each term separately : (a) $r_1 = r_2 = r_3 = 0$ (b) (r_1, r_2, r_3) take values $(0, 1, -1)$ in some order (c) $r_1 = r_2 = r_3 = 1$ (d) $r_1 = r_2 = r_3 = -1$.

The case (a) is trivial, the expectation being 1 in this case.

In case (b), lets say $(r_1, r_2, r_3) = (1, -1, 0)$, the other cases being similar. The expectation is

$$E_{w_1,w_2,f,g}[A(f \circ \pi)\overline{B(g \circ \pi')}]$$

Since π, π' are bijections, $f \circ \pi$ and $g \circ \pi'$ are distributed identically as f and g respectively. Hence the expectation is

$$E_{w_1,w_2,f,g}[A(f)\overline{B(g)}] = E_{w_1,f}[A(f)] \cdot \overline{E_{w_2,g}[B(g)]}$$

For a fixed v, let $\theta = E_{w,f}[A(f)] = E_w[\widehat{A}_0]$. Since w_1, w_2 are identically distributed, the above expectation is same as

$$E_{w,f}[A(f)] \cdot \overline{E_{w,f}[A(f)]} = |E_{w,f}[A(f)]|^2 = |\theta|^2$$

Now consider case (c). The expectation is

$$E[A(f \circ \pi)B(g \circ \pi')C((\overline{f}\overline{g}) \circ \pi'' \cdot \mu)]$$

Substituting Fourier expansions of A, B, C, we get

$$E\Big[\sum_{\alpha,\beta,\gamma} \widehat{A}_\alpha \widehat{B}_\beta \widehat{C}_\gamma \cdot \chi_\alpha(f \circ \pi)\ \chi_\beta(g \circ \pi')\ \chi_\gamma((\overline{f}\overline{g}) \circ \pi'' \cdot \mu) \Big]$$

Note that

$$\chi_\alpha(f \circ \pi) = \prod_{x \in M} f(\pi(x))^{\alpha(x)} = \prod_{y \in M} f(y)^{\alpha(\pi^{-1}(y))} = \chi_{\pi(\alpha)}(f)$$

where we define, by an abuse of notation, $\pi(\alpha)$ to be the function $\alpha \circ \pi^{-1}$. The previous expression reduces to

$$E\Big[\sum_{\alpha,\beta,\gamma} \widehat{A}_\alpha \widehat{B}_\beta \widehat{C}_\gamma \cdot \chi_{\pi(\alpha)-\pi''(\gamma)}(f)\ \chi_{\pi'(\beta)-\pi''(\gamma)}(g)\chi_\gamma(\mu)\Big]$$

Taking expectation over f, g, we see that the terms in this summation are zero unless $\pi(\alpha) = \pi'(\beta) = \pi''(\gamma)$. Also it is easy to check that $E_\mu[\chi_{\pi(\gamma)}(\mu)] = (\frac{-1}{2})^{|\gamma|}$ where for a function $\gamma : M \to GF(3)$ we define $|\gamma|$ to be the number of $x \in M$ such that $\gamma(x) \neq 0$. Thus the expectation reduces to

$$E_{w_1,w_2,w_3}\Big[\sum_{\pi(\alpha)=\pi'(\beta)=\pi''(\gamma)} \widehat{A}_\alpha \widehat{B}_\beta \widehat{C}_\gamma (-\frac{1}{2})^{|\gamma|} \Big] \qquad (5)$$

We will show that if the terms with $\gamma \neq 0$ are not small, one can extract labels for the Label Cover instance $\mathcal{L}$ giving a "good" value of $OPT(\mathcal{L})$. Lets assume

$$\delta \leq \left| E_{v,w_1,w_2,w_3}\Big[\sum_{\pi(\alpha)=\pi'(\beta)=\pi''(\gamma)\neq 0} \widehat{A}_\alpha \widehat{B}_\beta \widehat{C}_\gamma (-\frac{1}{2})^{|\gamma|} \Big] \right|$$

Applying Cauchy-Schwartz, this expression can be bounded by

$$E_{v,w_1,w_2,w_3}\Big[\sqrt{\sum_\alpha |\widehat{A}_\alpha|^2} \sqrt{\sum_{\pi'(\beta)=\pi''(\gamma)\neq 0} |\widehat{B}_\beta|^2 |\widehat{C}_\gamma|^2 (\frac{1}{4})^{|\gamma|}} \Big]$$

implying that

$$\delta^2 \leq E_{v,w_2,w_3}\Big[\sum_{\pi'(\beta)=\pi''(\gamma)\neq 0} |\widehat{B}_\beta|^2 |\widehat{C}_\gamma|^2 \frac{1}{|\gamma|} \Big]$$

Now we can define labels as follows. For a vertex $w_3 \in W$, pick γ with probability $|\widehat{C}_\gamma|^2$, pick a random $y \in M$ with $\gamma(y) \neq 0$ and define it to be the label of w_3. For a vertex $v \in V$, pick a random $w_2 \in \Psi_v$, pick β with probability $|\widehat{B}_\beta|^2$, pick a random $x \in M$ with $\beta(x) \neq 0$ and define $\pi'(x)$ to be the label of v. It is easy to see that this gives a labelling with $OPT(\mathcal{L}) \geq \delta^2$.

Hence we can choose $\mathcal{L}$ such that $OPT(\mathcal{L})$ is sufficiently small and ensure that the terms with $\gamma \neq 0$ in (5) are arbitrarily small. The term with $\gamma = 0$ contributes

$$E_{w_1,w_2,w_3}[\widehat{A}_0 \widehat{B}_0 \widehat{C}_0] \ = (E_w[\widehat{A}_0])^3 = \theta^3$$

The case (d) is just the complex conjugate of case (c) and it contributes $\overline{\theta}^3$ and terms which can be assumed to be arbitrarily small. Thus we can write the acceptance probability as

$$\Pr[Acc] = 1 - \frac{1}{9} - \frac{6}{9} E_v[|\theta|^2] - \frac{1}{9} E_v[\theta^3 + \overline{\theta}^3]$$

$$+ \text{ (Terms with arbitrarily small magnitude)}$$

Since $|\theta| \leq 1$, this probability is maximized when $\theta = 0$ (the reader familiar with this area will recognize that this means the proof better contain folded tables. If tables are not folded, it can only decrease the acceptance probability). Hence acceptance probability can be bounded by $\frac{8}{9} + \eta$ for arbitrarily small $\eta > 0$.

4.2 Hardness of 3-uniform Hypergraph Semi-coloring

We show that the Unique Games Conjecture implies that it is NP-hard to semi-color a 3-uniform hypergraph with constantly many colors when the hypergraph is given to be semi-colorable with 3 colors. This is proved by combining the techniques in the previous section with the idea of *covering complexity* of PCPs introduced by Guruswami et al [13]. We skip the proof.

5. PROOF OF THEOREM 1

In this section we prove Theorem 1. Instead of unique 2-prover games, we work in a more general setting of constraint satisfaction problems with uniqueness property.

Problem : *We are given a set X of n variables which take values from the set $[k] = \{1, 2, \ldots, k\}$. For every pair*

(u,v) of variables, there is a "constraint" which is a bijection $\pi_{uv} : [k] \to [k]$. This constraint has a weight w_{uv} with $\sum_{(u,v)} w_{uv} = 1$.

For an assignment $\mathcal{A} : X \to [k]$ to the variables, a constraint on the pair (u,v) is satisfied, if $\pi_{uv}(\mathcal{A}(u)) = \mathcal{A}(v)$. The goal is to find an assignment that maximizes the total weight of satisfied constraints.

Algorithm : We use a semidefinite program from Feige and Lovasz's paper [10] and augment it with a suitable rounding procedure. Let us first formulate the problem as a quadratic integer program. For every variable $u \in X$, let $u_1, u_2, \ldots, u_k$ be auxiliary variables taking 0-1 values. Place the following constraints :

$$u_1^2 + u_2^2 + \ldots u_k^2 = 1 \quad \forall\ u \in X \tag{6}$$

$$u_i u_j = 0 \quad \forall\ u \in X \text{ and } \forall\ i \neq j \tag{7}$$

We intend that if an assignment assigns the value $i_0 \in [k]$ to a variable u, then $u_{i_0} = 1$ and $u_i = 0\ \forall i \neq i_0$. This would satisfy the constraints (6), (7). These constraints imply that for every pair (u,v) of variables

$$u_i v_j \geq 0 \quad \forall\ i,j \tag{8}$$

$$\sum_{1 \leq i,j \leq k} u_i v_j = 1 \tag{9}$$

It is easy to see that the goal is to maximize the following function subjected to the above constraints.

$$\sum_{(u,v)} w_{uv}(u_1 v_{\pi(1)} + u_2 v_{\pi(2)} + \ldots u_k v_{\pi(k)}) \quad \text{where} \quad \pi = \pi_{uv} \tag{10}$$

Now we consider the semidefinite programming relaxation of the problem. We allow the variables $(u_1, \ldots, u_k)$ to be vectors in a high dimensional space (in kn-dimensional space to be precise) and the constraints (6)-(9) replaced by the constraints :

$$\vec{u}_1 \cdot \vec{u}_1 + \vec{u}_2 \cdot \vec{u}_2 + \ldots + \vec{u}_k \cdot \vec{u}_k = 1 \quad \forall\ u \in X \tag{11}$$

$$\vec{u}_i \cdot \vec{u}_j = 0 \quad \forall\ u \in X\ \ \forall\ \ i \neq j \tag{12}$$

$$\vec{u}_i \cdot \vec{v}_j \geq 0 \quad \forall\ u,v \in X\ \ \forall\ \ i,j \tag{13}$$

$$\sum_{1 \leq i,j \leq k} \vec{u}_i \cdot \vec{v}_j = 1 \quad \forall\ u,v \in X \tag{14}$$

The goal is to maximize the following function subjected to the above constraints :

$$\sum_{(u,v)} w_{uv}(\vec{u}_1 \cdot \vec{v}_{\pi(1)} + \ldots + \vec{u}_k \cdot \vec{v}_{\pi(k)}) \text{ where } \pi = \pi_{uv} \tag{15}$$

Observation : In any feasible solution of the SDP, for any two variables u, v, we have from the constraints (11), (12) and (14),

$$\|\sum_{i=1}^{k} \vec{u}_i\| = \|\sum_{j=1}^{k} \vec{v}_j\| = 1 \qquad \text{and} \qquad (\sum_{i=1}^{k} \vec{u}_i) \cdot (\sum_{j=1}^{k} \vec{v}_j) = 1$$

This implies that $\sum_{i=1}^{k} \vec{u}_i = \sum_{j=1}^{k} \vec{v}_j$. We denote $\vec{s} = \sum_{i=1}^{k} \vec{u}_i$ which is the same for all variables u and $\|\vec{s}\| = 1$.

We solve the semidefinite program and construct an assignment using the following rounding procedure.

- Choose a vector $\vec{r}$ from the normal distribution, i.e. choose every coordinate of $\vec{r}$ from the distribution $N(0,1)$ independently.
- By replacing $\vec{r}$ by $-\vec{r}$ if needed, assume that $\vec{r} \cdot \vec{s} \geq 0$.
- Construct the following assignment $\mathcal{A}$: for every variable u, let

$$\mathcal{A}(u) = i_0 \text{ where } \quad \vec{r} \cdot \vec{u}_{i_0} = \max_{1 \leq i \leq k} (\vec{r} \cdot \vec{u}_i)$$

We prove the following theorem in Appendix A which is sufficient to prove Theorem 1.

THEOREM 6. *If there exists an assignment that satisfies constraints with total weight $1 - \epsilon$, then the above algorithm produces an assignment that satisfies constraints with expected weight $1 - O(k^2 \epsilon^{1/5} \sqrt{\log(\frac{1}{\epsilon})})$.*

6. CONCLUSION

It seems quite difficult to prove (or disprove) the Unique Games Conjecture.

Proving the conjecture is equivalent to constructing a PCP that reads 2 symbols and accepts iff these symbols satisfy a bijective constraint. However the current tools appear quite weak for constructing PCPs that read 2 symbols. Parallel repetition of a unique game is a unique game and one might hope to amplify the soundness by parallel repetition. However we do not have a hard instance of a unique game to begin with. Theorem 1 shows that if the Unique Games Conjecture is true, the domain size $k(\zeta, \delta) \geq \frac{1}{\zeta^{1/10}}$, thus the domain size would play a very crucial role.

On the other hand, disproving the conjecture may require an algorithm that gives a theorem similar to Theorem 1 and whose performance is independent of the domain size k.

A less ambitious goal (than proving the Unique Games Conjecture) would be to show that the value of a unique 2-prover game with domain size k is hard to approximate within factor $f(k)$ where $f(k) \to \infty$ as $k \to \infty$. The only known results are constant factor hardness for 2-Linear-Equations mod 2 by Håstad [15] and for 2-Linear-Equations mod p by Andersson et al [2].

7. ACKNOWLEDGEMENT

I am greatful to Johan Håstad for showing me his analysis using Bourgain's theorem, which led me think about unique games. I thank Sanjeev Arora, Venkatesan Guruswami and Johan Håstad for many helpful discussions and their valuable comments on an earlier version of this paper.

8. REFERENCES

[1] N. Alon and N. Kahale. Approximating the independence number via the θ-function. *Technical Report, Tel Aviv University*, 1995.

[2] G. Andersson, L. Engebretsen, and J. Håstad. A new way of using semidefinite programming with applications to linear equations mod p. *Journal of Algorithms*, 39(2):162–204, 2001.

[3] S. Arora, C. Lund, R. Motawani, M. Sudan, and M. Szegedy. Proof verification and the hardness of approximation problems. *Journal of the ACM*, 45(3):501–555, 1998.

[4] S. Arora and S. Safra. Probabilistic checking of proofs : A new characterization of np. *Journal of the ACM*, 45(1):70–122, 1998.

[5] M. Bellare, O. Goldreich, and M. Sudan. Free bits, pcps and non-approximability. *Electronic Colloquium on Computational Complexity, Technical Report TR95-024*, 1995.

[6] J. Bourgain. On the distribution of the fourier spectrum of boolean functions. *manuscript*.

[7] I. Dinur and S. Safra. The importance of being biased. In *Proc. of the 34thth Annual ACM Symposium on Theory of Computing*, 2002.

[8] U. Feige. Error reduction - the state of the art. *Technical Report CS95-32, Weizmann Institute of Technology*, 1995.

[9] U. Feige. A threshold of ln n for approximating set cover. *Journal of the ACM*, 45(4):634–652, 1998.

[10] U. Feige and L. Lovasz. Two-prover one-round proof systems, their power and their problems. In *Proc. of the 24th Annual ACM Symposium on Theory of Computing*, pages 733–744, 1992.

[11] A. Frieze and M. Jerrum. Improved approximation algorithms for max k-cut and max bisection. *Algorihmica*, 18:67–81, 1997.

[12] M. Goemans and D. Williamson. 0.878 approximation algorithms for max-cut and max-2sat. In *Proc. of the 26th Annual ACM Symposium on Theory of Computing*, pages 422–431, 1994.

[13] V. Guruswami, J. Håstad, and M. Sudan. Hardness of approximate hypergraph coloring. In *Proc. of the 41st IEEE Symposium on Foundations of Computer Science*, pages 149–158, 2000.

[14] J. Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Proc. of the 37th Annual IEEE Symposium on Foundations of Computer Science*, pages 627–636, 1996.

[15] J. H. stad. Some optimal inapproximability results. In *Proc. of the 29th Annual ACM Symposium on Theory of Computing*, pages 1–10, 1997.

[16] J. Håstad. On a protocol possibly useful for min-2sat. *unpublished manuscript*.

[17] V. Kann, S. Khanna, J. Lagergren, and A. Panconesi. On the hardness of max k-cut and its dual. In *Proc. of the 5th Israel Symposium on Theory and Computing Systems*, pages 61–67, 1996.

[18] S. Khot. Hardness rsults for approximate hypergraph coloring. In *Proc. of the 34thth Annual ACM Symposium on Theory of Computing*, 2002.

[19] P. Klein, S. Plotkin, S. Rao, and E. Tardos. Approximation algorithms for steiner and directed multicuts. *Journal of Algorithms*, 22(2):241–269, 1997.

[20] R. Raz. A parallel repetition theorem. *SIAM J. of Computing*, 27(3):763–803, 1998.

[21] A. Samorodnitsky and L. Trevisan. A pcp characterization of np with optimal amortized query complexity. In *Proc. of the 32nd Annual ACM Symposium on Theory of Computing*, pages 191–199, 2000.

[22] U. Zwick. Approximation algorithms for constraint satisfaction problems involving at most three variables per constraint. In *Proc. of the 9th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 201–210, 1998.

[23] U. Zwick. Finding almost satisfying assignments. In *Proc. of the 30th Annual ACM Symposium on Theory of Computing*, pages 551–560, 1998.

APPENDIX

A. PROOF OF THEOREM 6

Let $\alpha_{uv} = \sum_{1 \le i \le k} \vec{u}_i \cdot \vec{v}_{\pi(i)},\ \pi = \pi_{uv}$ which is the part of the SDP objective function (15) corresponding to the constraint on (u, v). By the hypothesis, the SDP has a solution with value at least $1 - \epsilon$ implying that there exist vectors $(\vec{u}_i)_{u \in X, i \in [k]}$ satisfying

$$\sum_{(u,v)} w_{uv} \alpha_{uv} \ge 1 - \epsilon$$

$$\Longrightarrow \sum_{\alpha_{uv} \ge 1 - \frac{1}{2}\epsilon^{4/5}} w_{uv} \ge 1 - 2\epsilon^{1/5}$$

Fix any (u, v) with $\alpha_{uv} \ge 1 - \frac{1}{2}\epsilon^{4/5}$. We will show that with probability $1 - O(k^2 \epsilon^{1/5} \sqrt{\log(\frac{1}{\epsilon})})$, $\pi_{uv}(\mathcal{A}(u)) = \mathcal{A}(v)$. Let $\pi = \pi_{uv}$ for simplicity. The intuition behind the proof is simple : if $\alpha_{uv} = 1$, the SDP constraints (11-14) imply that $\vec{u}_i = \vec{v}_{\pi(i)}\ \forall\ i \in [k]$ (this can be seen by substituting $\epsilon = 0$ in Lemma 2). Thus for any vector $\vec{r}$, if $\vec{r} \cdot \vec{u}_i$ is maximized for index i_0, then $\vec{r} \cdot \vec{v}_j$ is maximized at index $\pi(i_0)$. Hence the rounding procedure will assign, $\mathcal{A}(u) = i_0$, and $\mathcal{A}(v) = \pi(i_0)$ satisfying the constraint.

We however have $\alpha_{uv} \ge 1 - \frac{1}{2}\epsilon^{4/5}$ and it takes some effort to translate the intuition into a rigorous proof. We proceed to prove several simple lemmas.

LEMMA 2. $\|\vec{u}_i - \vec{v}_{\pi(i)}\| \le \epsilon^{2/5} \quad \forall\ i \in [k]$.

PROOF.

$$1 - \frac{1}{2}\epsilon^{4/5} \le \sum_i \vec{u}_i \cdot \vec{v}_{\pi(i)} \le \sum_i \|\vec{u}_i\| \|\vec{v}_{\pi(i)}\|$$

$$\le \sum_i \frac{\|\vec{u}_i\|^2 + \|\vec{v}_{\pi(i)}\|^2}{2} = 1$$

$$\Longrightarrow \frac{\|\vec{u}_i\|^2 + \|\vec{v}_{\pi(i)}\|^2}{2} - \vec{u}_i \cdot \vec{v}_{\pi(i)} \le \frac{1}{2}\epsilon^{4/5} \quad \forall\ i$$

$$\Longrightarrow \|\vec{u}_i - \vec{v}_{\pi(i)}\|^2 \le \epsilon^{4/5} \quad \forall\ i$$

□

LEMMA 3. *If Y is distributed as $N(0,1)$,*

$$Pr\left[|Y| > \gamma\right] \le e^{\frac{-\gamma^2}{2}}$$

PROOF. Standard inequality. □

LEMMA 4. *With probability $1 - O(k^2 \epsilon^{1/5} \sqrt{\log(\frac{1}{\epsilon})})$, components of $\vec{r}$ along the directions of vectors*

$$\{\vec{u}_i\}_{i \in [k]}, \{\vec{u}_i - \vec{u}_j\}_{i \ne j},\ \{\vec{u}_i - \vec{v}_{\pi(i)}\}_{i \in [k]}$$

have magnitude in the range

$$\left[\ \epsilon^{1/5} \sqrt{\log(\frac{1}{\epsilon})},\ \sqrt{\log(\frac{1}{\epsilon})}\ \right]$$

PROOF. This follows from the fact that $\vec{r}$ is distributed in a spherically symmetric manner and hence its component along any direction is distributed as $N(0,1)$. Hence for any unit vector $\vec{t}$,

$$\Pr\left[\ |\vec{r}\cdot\vec{t}| < \epsilon^{1/5}\sqrt{\log(\frac{1}{\epsilon})}\ \right] \ < \ 2\epsilon^{1/5}\sqrt{\log(\frac{1}{\epsilon})}$$

$$\Pr\left[\ |\vec{r}\cdot\vec{t}| > \sqrt{\log(\frac{1}{\epsilon})}\ \right] \ < \ \sqrt{\epsilon}$$

where the first inequality is trivial and the second follows from Lemma 3. Now we take a union bound along the $O(k^2)$ directions specified in the statement of this lemma. □

LEMMA 5. *With probability* $1-10k\epsilon^{1/5}\sqrt{\log(\frac{1}{\epsilon})}$, *the component of* $\vec{r}$ *along* $\vec{s}$, *that is* $|\vec{r}\cdot\vec{s}|$, *is at least* $5k\epsilon^{1/5}\sqrt{\log(\frac{1}{\epsilon})}$.

PROOF. Trivial. □

Thus except with probability $1-O(k^2\epsilon^{1/5}\sqrt{\log(\frac{1}{\epsilon})})$, we can assume that $\vec{r}$ satisfies hypothesis of Lemma 4 and Lemma 5. Under this assumption, we prove the following 3 lemmas. Let $i_0 \in [k]$ be such that $\vec{r}\cdot\vec{u}_{i_0} = \max_{1\le i\le k}\vec{r}\cdot\vec{u}_i$.

LEMMA 6. $\|\vec{u}_{i_0}\| \ge 5\epsilon^{1/5}$.

PROOF. $(\sum_{i=1}^{k}\vec{u_i})\cdot\vec{r} = \vec{s}\cdot\vec{r} \ge 5k\epsilon^{1/5}\sqrt{\log(\frac{1}{\epsilon})}$ by Lemma 5 and i_0 is the index that maximizes $\vec{r}\cdot\vec{u}_i$. Hence $\vec{r}\cdot\vec{u}_{i_0} \ge 5\epsilon^{1/5}\sqrt{\log(\frac{1}{\epsilon})}$. But by Lemma 4, the component of $\vec{r}$ along $\vec{u}_{i_0}$ has magnitude at most $\sqrt{\log(\frac{1}{\epsilon})}$. This implies that $\|\vec{u}_{i_0}\| \ge 5\epsilon^{1/5}$. □

LEMMA 7. $\forall\ j \ne i_0,\quad \vec{r}\cdot\vec{u}_j \le \vec{r}\cdot\vec{u}_{i_0} - 5\epsilon^{2/5}\sqrt{\log(\frac{1}{\epsilon})}$

PROOF.

$$\begin{aligned}
\vec{r}\cdot\vec{u}_{i_0} - \vec{r}\cdot\vec{u}_j &= |\vec{r}\cdot\vec{u}_{i_0} - \vec{r}\cdot\vec{u}_j| \\
&= |\vec{r}\cdot(\vec{u}_{i_0} - \vec{u}_j)| \\
&\ge \|\vec{u}_{i_0} - \vec{u}_j\|\ \ \epsilon^{1/5}\sqrt{\log(\tfrac{1}{\epsilon})} \quad \text{by Lemma 4} \\
&\ge \|\vec{u}_{i_0}\|\ \epsilon^{1/5}\sqrt{\log(\tfrac{1}{\epsilon})} \qquad \text{Since } \vec{u}_{i_0}\perp\vec{u}_j \\
&\ge 5\epsilon^{2/5}\sqrt{\log(\tfrac{1}{\epsilon})} \qquad \text{by Lemma 6}
\end{aligned}$$

□

LEMMA 8. $\forall\ i,\quad |\vec{r}\cdot\vec{u}_i - \vec{r}\cdot\vec{v}_{\pi(i)}| \le \epsilon^{2/5}\sqrt{\log(\frac{1}{\epsilon})}$

PROOF.

$$\begin{aligned}
|\vec{r}\cdot\vec{u}_i - \vec{r}\cdot\vec{v}_{\pi(i)}| &= |\vec{r}\cdot(\vec{u}_i - \vec{v}_{\pi(i)})| \\
&\le \sqrt{\log(\tfrac{1}{\epsilon})}\ \|\vec{u}_i - \vec{v}_{\pi(i)}\| \quad \text{by Lemma 4} \\
&\le \sqrt{\log(\tfrac{1}{\epsilon})}\epsilon^{2/5} \qquad \text{by Lemma 2}
\end{aligned}$$

□

Now we will show that

$$\vec{r}\cdot\vec{v}_{\pi(i_0)} = \max_{1\le j\le k}(\vec{r}\cdot\vec{v_j}) \tag{16}$$

This would imply that the assignment $\mathcal{A}$ given by the rounding procedure assigns $\mathcal{A}(u) = i_0,\ \mathcal{A}(v) = \pi(i_0)$ and the constraint on the pair (u,v) is satisfied.

Let $j \ne i_0$ be any index. By Lemma 8 and Lemma 7,

$$\vec{r}\cdot\vec{v}_{\pi(j)} \le \vec{r}\cdot\vec{u}_j + \epsilon^{2/5}\sqrt{\log(\frac{1}{\epsilon})} \le \vec{r}\cdot\vec{u}_{i_0} - 4\epsilon^{2/5}\sqrt{\log(\frac{1}{\epsilon})}$$

Also by Lemma (8) we have

$$\vec{r}\cdot\vec{v}_{\pi(i_0)} \ge \vec{r}\cdot\vec{u}_{i_0} - \epsilon^{2/5}\sqrt{\log(\frac{1}{\epsilon})}$$

It follows that

$$\vec{r}\cdot\vec{v}_{\pi(i_0)} > \vec{r}\cdot\vec{v}_{\pi(j)} \quad \forall\ j \ne i_0$$

finishing the proof of (16) and Theorem 6.