



Improved 3LIN Hardness via Linear Label Cover

Prahladh Harsha^{*} Subhash Khot[†] Euiwoong Lee[†]
 Devanathan Thiruvengatachari[†]

Abstract

We prove that for every constant c and $\varepsilon = (\log n)^{-c}$, there is no polynomial time algorithm that when given an instance of 3-LIN with n variables where an $(1 - \varepsilon)$ -fraction of the clauses are satisfiable, finds an assignment that satisfies at least $(\frac{1}{2} + \varepsilon)$ -fraction of clauses unless $\mathbf{NP} \subseteq \mathbf{BPP}$. The previous best hardness using a *polynomial time* reduction achieves $\varepsilon = (\log \log n)^{-c}$, which is obtained by the LABEL COVER hardness of Moshkovitz and Raz [*J. ACM*, 57(5), 2010] followed by the reduction from LABEL COVER to 3-LIN of Håstad [*J. ACM*, 48(4):798–859, 2001].

Our main idea is to prove a hardness result for LABEL COVER similar to Moshkovitz and Raz where each projection has a *linear* structure. This linear structure of LABEL COVER allows us to use Hadamard codes instead of long codes, making the reduction more efficient. For the hardness of LINEAR LABEL COVER, we follow the work of Dinur and Harsha [*SIAM J. Comput.*, 42(6):2452–2486, 2013] that simplified the construction of Moshkovitz and Raz, and observe that running their reduction from a hardness of the problem LIN (of unbounded arity) instead of the more standard problem of solving quadratic equations ensures the linearity of the resultant LABEL COVER.

1 Introduction

In this paper, we study the 3-LIN problem. An instance of 3-LIN consists of a set of n variables over \mathbb{F}_2 and a set of m equations that contain at most three variables each, and the goal is to find an assignment to the n variables that satisfies the most number of equations.¹ If the given set of linear equations admits an assignment that satisfies every equation, then one such assignment can be found in polynomial time by Gaussian elimination. However, the general problem of finding the most of number of equations is \mathbf{NP} -hard when the instance does not admit a satisfying assignment, and a large amount of research has been done on the limit of polynomial time approximation algorithms.

Assigning random values to variables satisfies exactly half the equations in expectation, giving a $1/2$ -approximation algorithm. Håstad and Venkatesh [HV04] achieved an approximation factor of $1/2 + 1/O(\sqrt{m})$, which was improved by Khot and Naor [KN08] to $1/2 + O(\sqrt{\log n/n})$.

^{*}School of Technology and Computer Science, Tata Institute of Fundamental Research, India. Supported in part by the DIMACS/Simons Collaboration in Cryptography through NSF grant #CNS-1523467 (while the author was visiting Rutgers University and DIMACS) and the Swarnajayanti Fellowship.

[†]Department of Computer Science, Courant Institute of Mathematical Sciences, New York University, USA. Supported by the NSF Award CCF-1422159, the Simons Collaboration on Algorithms and Geometry and the Simons Investigator Award.

¹This maximization version is also known as MAX 3-LIN in the literature.

From the hardness side, there are strong hardness results even when the instance is *almost-satisfiable*. For $1 \geq c > s > 0$, let $\text{GAP 3-LIN}(c, s)$ denote the problem of distinguishing whether the given instance of 3-LIN is at least c -satisfiable or at most s -satisfiable. Håstad’s classic hardness results [Hås01] show the following.

Theorem 1.1 ([Hås01]). *The following hardness results for GAP 3-LIN hold.*

1. For any constant $\varepsilon > 0$, $\text{GAP 3-LIN}(1 - \varepsilon, 1/2 + \varepsilon)$ is **NP**-hard.
2. There exists a constant $c > 0$ such that for $\varepsilon = 1/(\log n)^c$, there is no polynomial time algorithm that solves $\text{GAP 3-LIN}(1 - \varepsilon, 1/2 + \varepsilon)$ unless $\mathbf{NP} \subseteq \mathbf{DTIME}[n^{O(\log \log n)}]$.

Håstad’s results are proved by giving the reduction from LABEL COVER to 3-LIN. LABEL COVER is a common starting point for hardness results, and we define the optimization problem below.

Definition 1.2 (LABEL COVER). *An instance of LABEL COVER contains a regular bipartite multi-graph $G = (A, B, E)$ and two finite sets Σ_A and Σ_B , where $|\Sigma_A| \geq |\Sigma_B|$. Every vertex in A is supposed to get a label in Σ_A , and every vertex in B is supposed to get a label in Σ_B . For each edge $e \in E$ there is a projection $\pi_e : \Sigma_A \rightarrow \Sigma_B$. Given a labeling to the vertices of the graph, i.e., functions $\phi_A : A \rightarrow \Sigma_A$ and $\phi_B : B \rightarrow \Sigma_B$, an edge $e = (a, b) \in E$ is said to be “satisfied” if $\pi_e(\phi_A(a)) = \phi_B(b)$. For $1 \geq c > s > 1$, $\text{GAP LABEL COVER}(c, s)$ is the problem of distinguishing whether the given instance of LABEL COVER is at least c -satisfiable or at most s -satisfiable.*

Håstad’s theorem can be stated in terms of reduction from $\text{GAP LABEL COVER}(1, \delta)$ as follows.

Theorem 1.3 ([Hås01]). *For every $\varepsilon \in (0, 1)$ and positive integer ℓ , there exists a $\delta = \text{poly}(\varepsilon)$ and a $\text{poly}(n, 2^\ell, 2^{1/\varepsilon})$ -time reduction from n -sized instances of $\text{GAP LABEL COVER}(1, \delta)$ with label size ℓ to $\text{GAP 3-LIN}(1 - \varepsilon, 1/2 + \varepsilon)$.*

When [Hås01] was published, the hardness of LABEL COVER was achieved by the PCP Theorem [AS98, ALM⁺98] and parallel repetition [Raz98]. More precisely, $\text{GAP LABEL COVER}(1, \varepsilon)$ with label size $\text{poly}(1/\delta)$ was **NP**-hard under $\text{poly}(n^{\log 1/\delta})$ -time reductions. The two results of Håstad stated in theorem 1.1 follow from this hardness of GAP LABEL COVER and theorem 1.3 by setting δ to be an arbitrarily small constant and $1/\log n$ respectively. Since achieving a subconstant soundness for LABEL COVER by parallel repetition requires a superpolynomial blowup in the instance size, $\varepsilon > 0$ could not be taken to subconstant under *polynomial time* reductions. Later in a celebrated paper, Moshkovitz and Raz [MR10] gave an improved hardness of LABEL COVER that achieves sub-constant error under polynomial time reductions. Their main result can be stated as follows.

Theorem 1.4 ([MR10, Theorem 11]). *For every n , and every $\delta > 0$ (that can be any function of n), 3-SAT on inputs of size n can be reduced to $\text{GAP LABEL COVER}(1, \delta)$ when LABEL COVER instance has $n^{1+o(1)} \cdot \text{poly}(1/\delta)$ vertices and $|\Sigma_A| \leq \exp(\text{poly}(1/\delta))$, $|\Sigma_B| \leq \text{poly}(\log 1/\delta)$.*

A corollary of the above result, obtained by combining it with Håstad’s reduction from theorem 1.3, is that given a system of linear equations, it is **NP**-hard to distinguish between cases where $1 - o(1)$ fraction of equations are satisfied vs at most $1/2 + o(1)$ fraction are satisfied, where the $o(1)$ term is $1/(\log \log n)^{-\Omega(1)}$.

Theorem 1.5 ([MR10]). *There exists some constant $c > 0$ such that for $\varepsilon = 1/(\log \log n)^c$, GAP 3-LIN($1 - \varepsilon, 1/2 + \varepsilon$) is **NP**-hard.*

Later, an improved parallel repetition by Dinur and Steurer [DS14] allowed c to be an arbitrary constant.

The above route prove hardness of 3-LIN is restricted by the large size of the alphabet in the resulting LABEL COVER instance in theorem 1.4. Quantitatively, the alphabet size is exponential in $\text{poly}(1/\varepsilon)$. The fact that the long code in Håstad's reduction has size exponential in the alphabet size restricts $\varepsilon = 1/(\log \log n)^{O(1)}$.

Our main contribution for 3-LIN is to bring ε in the above result down to $1/(\log n)^c$ for any constant c , while keeping the size of the reduced instance polynomial (albeit the reduction becomes randomized).

Theorem 1.6 (Main). *For any constant $c > 0$ and $\varepsilon = 1/(\log n)^c$, there is no polynomial time algorithm for GAP 3-LIN($1 - \varepsilon, 1/2 + \varepsilon$) unless **NP** \subseteq **BPP**.*

We get around the above alphabet barrier by starting with a reduction that would make the resulting LABEL COVER *linear*, and use Hadamard codes instead of long codes. Since the Hadamard code keeps the reduction size polynomial in the alphabet size, we can take $\varepsilon = 1/(\log n)^{\Omega(1)}$. A similar idea was previously used by Khot [Kho01]. We define LINEAR LABEL COVER as follows.

Definition 1.7 (LINEAR LABEL COVER). *A LINEAR LABEL COVER is a special case of LABEL COVER where the alphabets are of the form $\Sigma_A = \mathbb{F}_2^a, \Sigma_B = \mathbb{F}_2^b$ where a, b are natural numbers. Each projection $\pi : \mathbb{F}_2^a \rightarrow \mathbb{F}_2^b$ is affine in the sense that $\pi(x) = \alpha x + \beta$ for some $\alpha \in \mathbb{F}_2^{b \times a}, \beta \in \mathbb{F}_2^b$. For $1 \geq c > s > 0$, the GAP LINEAR LABEL COVER(c, s) is defined similarly to GAP LABEL COVER(c, s).*

We prove the following hardness result for LINEAR LABEL COVER, which may be of independent interest.

Theorem 1.8 (Hardness of Linear Label Cover). *For any constant $c > 0$, for $\delta = 1/(\log n)^c$, there is no polynomial time algorithm for GAP LINEAR LABEL COVER($1 - \delta, \delta$) unless **NP** \subseteq **BPP**, when LABEL COVER instance has $\text{poly}(n)$ vertices and $|\Sigma_A| = \text{poly}(n), |\Sigma_B| = \text{polylog}(n)$.*

We remark that if the above theorem can be further strengthened to obtain $\delta = 1/n^c$ (i.e., a linear version of the Sliding Scale conjecture), then this leads to near-optimal hardness of 3-LIN (i.e, GAP 3-LIN($1 - \varepsilon, 1/2 + \varepsilon$) is hard for $\varepsilon = 1/\text{poly}(n)$) [Mos15].

1.1 Proof Ideas

Our main technical contribution is theorem 1.8 for LINEAR LABEL COVER, essentially proving a linear analogue of the Moshkovitz-Raz PCP [MR10] followed by the Dinur-Steurer parallel repetition [DS14]. The proof is given through a long sequence of reductions. We split them in 3 major steps.

1. Interestingly, the starting point of our reduction is again the hardness of (not necessarily linear) LABEL COVER proved by Moshkovitz and Raz [MR10] augmented by Dinur and Steurer [DS14], proving **NP**-hardness of GAP LABEL COVER($1, 1/\log^c n$) for any $c > 0$, while keeping the reduction size and the alphabet size polynomial. In section 2, we give a *randomized*

reduction from this LABEL COVER to GAP LIN($1 - 1/\log^c n, 0.9$). This style of reduction appeared previous from LABEL COVER to CLOSEST VECTOR PROBLEM [Kho10]. Note that the standard proof of the PCP theorem encodes 3-SAT (or CIRCUIT SAT) by solving quadratic equations over \mathbb{F}_2 , and this is essentially the only place that needs where nonlinearity occurs. Our hardness result for solving linear equations with completeness very close to (but not exactly) 1 allows us to follow previous PCP constructions that will ensure linearity of the LABEL COVER instance in the subsequent steps.

2. To prove the hardness of LINEAR LABEL COVER given the above hardness of LIN, we closely follow the steps of Dinur and Harsha [DH13], who gave a simpler and modular proof of [MR10]. The two basic building blocks in their proof are robust PCPs and decodable PCPs. Robust PCPs are PCPs where in the soundness case, for any proof and most random choices of the verifier, not only are the local views non-accepting, but they are also very far from any accepting string. It is indeed equivalent to LABEL COVER. Using our previous hardness for LIN as the starting point and following the standard robust PCP construction (e.g., low-degree extension and sum-check protocol), we can prove a polynomial time reduction to LINEAR LABEL COVER($1 - 1/\log^c n, 1/\log^c n$) for any $c > 1$, but the alphabet size will be always $\exp(\log^{c_0} n)$ for some $c_0 > 1$, which is superpolynomial.
3. The second building block, decodable PCP, is similar to robust PCP with the additional requirement that the prover is given a position i in the original string and supposed to output the value of the i th position if the given proof is a honest encoding of a valid original string. The main idea of Dinur and Harsha [DH13] is to iteratively compose a robust PCP with a suitable decodable PCP, where the composed PCP is another robust PCP that consists of a decodable PCP for each constraint of the original robust PCP. This iteratively reduces the query complexity and the alphabet size of the robust PCP, which is related to the alphabet size of the equivalent LABEL COVER instance. This iterative composition is interleaved and preprocessed by technical operations that reduce the alphabet size of the robust PCP and make it regular.

Once these two building blocks are linear, the operations of [DH13] can be used verbatim in our construction. Our main observation is that every step of this construction preserves (1) the robust completeness $1 - \delta$ for some $\delta = 1/\text{polylog}(n)$, and (2) the linearity, which were not issues in [DH13]. In section 3, we introduce the basic building blocks and these operations, and show how they preserve robust completeness and linearity. These iterative operations will eventually reduce the alphabet size of the LINEAR LABEL COVER polynomial, proving theorem 1.8.

After the hardness of LINEAR LABEL COVER is proved, we give a reduction from LINEAR LABEL COVER with the above parameters to 3-LIN with the required parameters. We do this by composing with the Hadamard Code to get a $(1 - \varepsilon)$ vs $(1/2 + \varepsilon)$ NP-hardness result for 3LIN. Similar PCP constructions based on Hadamard codes were presented in [Kho01]. Details of this step can be found in section 4.

2 Reduction to System of Linear Equations

In this section, we first prove the hardness of approximate solving linear equations over large fields, where each equation can involve as many variables as possible. It will serve as the starting point towards proving hardness of LINEAR LABEL COVER.

Theorem 2.1. *For any constant $c > 0$, $\varepsilon = 1/(\log n)^c$, GAP LIN($1 - 1/(\log n)^c, 0.9$) is **NP**-hard under polynomial time randomized reductions.*

Proof. The proof starts from the following hardness of LABEL COVER, which is obtained by combining the main result of Moshkovitz and Raz [MR10] with the parallel repetition of Dinur and Steurer [DS14].

Theorem 2.2 ([MR10, DS14]). *For any constant $c > 0$, for $\delta = 1/(\log n)^c$, GAP LABEL COVER($1, \delta$) is **NP**-hard when the LABEL COVER instance satisfies $|\Sigma_A|, |\Sigma_B| \leq |A| + |B|$.*

Let $G = (A, B, E)$, Σ_A, Σ_B , and $\{\pi_e\}_{e \in E}$ be an instance of LABEL COVER. We show a reduction to LIN over \mathbb{F}_2 where

- If all LABEL COVER edges are satisfiable, at least $(1 - \frac{1}{|\Sigma_A|})$ fraction of equations are satisfiable.
- If at most δ fraction of LABEL COVER edges are satisfiable, at most $(1 - \frac{1}{\delta|\Sigma_A|})$ fraction of equations are satisfiable.

For each vertex $v \in \Sigma_A \cup \Sigma_B$ and possible label ℓ on the Label Cover instance, we have a variable $x_{v,\ell}$ in the LIN instance. Let $n = |A||\Sigma_A| + |B||\Sigma_B| = \text{poly}(|A| + |B|)$ be the number of variables. Consider the following four kinds of equations. Recall that every arithmetic is performed over \mathbb{F}_2 .

$$\begin{aligned}
 (1) \quad & \sum_{\ell \in \Sigma_A} x_{v,\ell} = 1 && \forall v \in A \\
 (2) \quad & \sum_{\ell \in \Sigma_B} x_{v,\ell} = 1 && \forall v \in B \\
 (3) \quad & \sum_{r: \pi_{uv}(r)=\ell} x_{v,r} = x_{u,\ell} && \forall (u, v) \in E, \forall \ell \in \Sigma_B \\
 (4) \quad & x_{v,\ell} = 0 && \forall (v, \ell) \in A \times \Sigma_A
 \end{aligned}$$

In our final LIN instance, we treat (1), (2), and (3) as *hard constraints* that need to be always satisfied, and find x that always satisfies all hard constraints and as many constraints in (4) as possible. Also note that in (4), we only consider vertices in A .

This is equivalent to the usual LIN problem with hard constraints by *folding*. Formally, let V be the set of assignments that satisfy (1), (2), and (3). If V is empty, we can conclude that the LABEL COVER instance is unsatisfiable. Otherwise, there exist $c \in \mathbb{N}$ and linearly independent vectors $y_0, \dots, y_c \in \mathbb{F}_2^{(A \times \Sigma_A) \cup (B \times \Sigma_B)}$ such that $V = \{y_0 + \sum_{i=1}^c y_i z_i : z_1, \dots, z_c \in \mathbb{F}_2\}$. This gives an one-to-one correspondence between \mathbb{F}_2^c and V , so we can treat z_1, \dots, z_c as the variables of LIN and write the fourth constraints $x_{v,\ell} = 0$ in terms of z , which gives an instance of LIN without hard constraints.

Completeness. If the LABEL COVER instance is satisfiable, $x_{v,\ell} = 1$ if and only if v is assigned with ℓ gives an assignment that satisfies (1), (2), and (3), and violates one equation in (4) for each $v \in A$.

Soundness. Let x be an assignment that satisfies (1), (2), and (3). For $v \in A \cup B$, let $L_v := \{\ell : x_{v,\ell} = 1\}$. Since (1) and (2) require $\sum_{\ell} x_{v,\ell} = 1$ for every $v \in A \cup B$, L_v is not empty for every v .

Consider the randomized strategy for LABEL COVER where each $v \in A \cup B$ is assigned with a uniform random label from L_v independently. For $(u, v) \in E$ with $u \in A, v \in B$, by (3), $x_{v,\ell} = 1$ for some $\ell \in \Sigma_B$ implies that there exists $r \in \Sigma_A$ with $\pi_{uv}(r) = \ell$ such that $x_{u,r} = 1$. This implies (u, v) is satisfied with probability at least $\frac{1}{|L_u|}$ by the randomized strategy. Then the expected fraction of the LABEL COVER constraints satisfied by the strategy is at least

$$\mathbf{E}_{u \in A} \left[\frac{1}{|L_u|} \right] \geq \frac{1}{\mathbf{E}_{u \in A}[|L_u|]}.$$

Therefore, if at most δ fraction of LABEL COVER constraints are simultaneously satisfiable, we can conclude that

$$\delta \geq \frac{1}{\mathbf{E}_{u \in A}[|L_u|]} \Leftrightarrow \mathbf{E}_{u \in A}[|L_u|] \geq \frac{1}{\delta}.$$

So in total, at least $\frac{1}{(\delta|\Sigma_A|)}$ fraction of equations are violated.

Gap Amplification. We have a hardness of LIN over \mathbb{F}_2 where the completeness value is at least $1 - \frac{1}{|\Sigma_A|}$ and the soundness value is at most $1 - \frac{1}{(\delta|\Sigma_A|)}$. Consider a new system of linear equations where we sample m linear equations independently, where each new equation randomly chooses $\delta \cdot |\Sigma_A|$ old equations and takes a random linear combination of them. In the completeness case, at least an $(1 - O(\delta))$ fraction of new equations can be satisfied by a good assignment to old equations.

In the soundness case, fix an assignment to n possible variables. (There are 2^n of them.) It satisfies at most an $1 - \frac{1}{(\delta|\Sigma_A|)}$ fraction of old equations. Note that if a new equation chooses an old equation not satisfied by the assignment, it is satisfied with probability exactly $1/2$. Therefore, the expected number of new equations satisfied by this fixed assignment is at most

$$m \cdot \left(\left(1 - \frac{1}{(\delta|\Sigma_A|)}\right)^{\delta \cdot |\Sigma_A|} + \frac{1}{2} \right) \leq m \cdot \left(\frac{1}{e} + \frac{1}{2} \right) \leq 0.87m.$$

For a given $c \in \mathbb{N}$, let $\delta = 1/\log^c n$. By taking sufficiently large $m = O(n)$, we can apply the Chernoff and union bound to conclude that no assignment satisfies more than a 0.9 fraction of new equations. So we reduce from LABEL COVER to GAP LIN($1 - O(\delta)$, 0.9), which finishes the proof. ■

We remark that the sampling performed above is the only step in our reduction involving randomization.

3 Reduction to Linear Label Cover

In this section, we show for any $c > 0$, unless $\mathbf{NP} \subseteq \mathbf{BPP}$, there is no polynomial time algorithm for GAP LINEAR LABEL COVER($1 - \varepsilon, \varepsilon$) with $\varepsilon = 1/(\log n)^c$, proving theorem 1.8.

The construction we employ is almost identical to that of Dinur and Harsha [DH13], except that the basic building blocks (robust PCP and decodable PCP) try to prove (almost) satisfiability of linear equations instead of standard quadratic equations. They are introduced in sections 3.1 and 3.2.

After constructing the building blocks, the result of [DH13] is proved by iterative composition of them followed by technical steps including alphabet and degree reduction. Our main observation in this part is that each of the steps in the construction preserves *linearity* so that the final LABEL COVER instance produced also has a linear structure. We present them in section 3.3 and section 3.4. Finally, section 3.5 shows how to combine all these steps to prove theorem 1.8.

3.1 Robust PCPs

In this subsection, we define robust PCPs. For two strings x, y of the same length n , let $\text{agr}(x, y)$ denote the relative agreement of the strings x, y , defined as

$$\text{agr}(x, y) := \Pr_{i \in [n]} [x_i = y_i]$$

If S is a set of strings, $\text{agr}(x, S)$ is defined as $\max_{y \in S} \{\text{agr}(x, y)\}$.

Definition 3.1 (Robust PCPs). *For functions $r, q, m, a, s : \mathbb{N} \rightarrow \mathbb{N}$ and $c, \delta : \mathbb{N} \rightarrow [0, 1]$, a verifier V is a robust probabilistically checkable proof (robust PCP) system for a promise problem $L = (L_{\text{YES}}, L_{\text{NO}})$ with randomness complexity r , query complexity q , proof length m , alphabet size a , robust completeness c , and robust soundness error δ if V is a probabilistic polynomial-time algorithm that behaves as follows: On input x of length n and oracle access to a proof string $\pi \in \Sigma^{m(n)}$ over the (proof) alphabet Σ where $|\Sigma| = a(n)$, V reads the input x , tosses at most $r = r(n)$ random coins, and generates a sequence of locations $I = (i_1, \dots, i_q) \in [m]^{q(n)}$ and a predicate $f : \Sigma^q \rightarrow \{0, 1\}$, which satisfy the following properties.*

Robust Completeness. *If $x \in L_{\text{YES}}$ then there exists π such that*

$$\mathbf{E}_{(I, f)} [\text{agr}(\pi_I, f^{-1}(1))] \geq c. \quad (1)$$

Robust Soundness. *If $x \in L_{\text{NO}}$ then for every π ,*

$$\mathbf{E}_{(I, f)} [\text{agr}(\pi_I, f^{-1}(1))] \leq \delta, \quad (2)$$

where the distribution over (I, f) is determined by x and the random coins of V .

We say that V is *linear* if $\Sigma = \mathbb{F}_2^b$ for some b and for every f , the accepting sets of the predicate f , i.e., $f^{-1}(1)$, forms an affine subspace of $\Sigma^q = \mathbb{F}_2^{bq}$ over the field \mathbb{F}_2 .

Robust completeness and soundness must be contrasted with (regular) completeness and soundness of standard PCP verifiers in which the expression for completeness and soundness given in (1) and (2) respectively are replaced as follows:

$$\text{Completeness: } \Pr_{I, f} [f(\pi_I) = 1] \geq c,$$

$$\text{Soundness: } \Pr_{I, f} [f(\pi_I) = 1] \leq \delta.$$

In fact, this is the only difference between the above definition and the standard definition of a PCP system. The robust soundness states that not only does the local view violate the local predicate f , but in fact has very little agreement with any of the satisfying assignments of f (and thus is a strengthening of standard robustness). Robust completeness on the other hand is a weakening of standard completeness.

Another crucial aspect of robust PCP is its equivalence to LABEL COVER. Namely, existence of robust PCP for L with parameters r, q, m, a, s, c, δ is equivalent to existence of a reduction from L to GAP LABEL COVER(c, δ) where $|A| = 2^r, |B| = m, |\Sigma_A| \leq a^q, |\Sigma_B| = a$ and each $v \in A$ has degree q . See Lemma 2.5 of [DH13]. Also note that the definition of linearity is equivalent in robust PCP and LABEL COVER.

Theorem 3.2 (Robust PCP, Analog of [DH13, Theorem 6.4]). *There exist constants $b_1, b_2 > 0, c_0 > 1$ such that for any $c > c_0$ and $\varepsilon = 1/\log^c n$, GAP LIN($1 - \varepsilon, 0.9$) with n variables has a linear robust verifier with robust completeness $1 - \varepsilon$, robust soundness error ε , query complexity $1/\varepsilon^{b_1}$, proof length $\text{poly}(n)$, randomness complexity $O(\log n)$, and proof alphabet size at most $1/\varepsilon^{b_2}$.*

Equivalently, there is a (deterministic) polynomial time reduction from GAP LIN($1 - \varepsilon, 0.9$) to GAP LINEAR LABEL COVER($1 - \varepsilon, \varepsilon$), where the LABEL COVER instance has $\text{poly}(n)$ veritces, $|\Sigma_A| \leq \exp(1/\varepsilon^{b_1} \log(1/\varepsilon^{b_2}))$, $|\Sigma_B| \leq 1/\varepsilon^{b_2}$, and each $v \in A$ has degree $1/\varepsilon^{b_1}$.

The proof of this theorem is identical to that of [DH13, Theorem 6.4] and omitted here. The only difference is GAP LIN($1 - \varepsilon, 0.9$) with $1/\varepsilon = \log^{O(c)} n$ instead of standard quadratic equations when performing the low degree-extension and the sum-check protocol. The theorem follows by observing that all the operations are linear and hence the final predicate is also linear. The completeness of the robust PCP is dictated by the completeness value in theorem 2.1.

Combining this reduction with the randomized reduction from theorem 2.1, we obtain the following theorem (which is a more formal version of theorem 1.8).

Theorem 3.3 (Hardness of Linear Label Cover). *There exist constants $b_1, b_2 > 0, c_0 > 1$ such that for any $c > c_0$ and $\varepsilon = 1/\log^c n$, unless $\mathbf{NP} \subseteq \mathbf{BPP}$, there is no polynomial time algorithm for GAP LINEAR LABEL COVER($1 - \varepsilon, \varepsilon$) where the LABEL COVER instance has $\text{poly}(n)$ veritces, $|\Sigma_A| \leq \exp(1/\varepsilon^{b_1} \log(1/\varepsilon^{b_2}))$, $|\Sigma_B| \leq 1/\varepsilon^{b_2}$, and each $v \in A$ has degree $1/\varepsilon^{b_1}$.*

3.2 Decodable PCPs

We now discuss the decodable PCP (dPCP), which differs from a PCP in that it has a decoder as opposed to a verifier. A *decoder* is similar to a verifier in that it checks whether a string is in the given language or not by probabilistically checking a small number of positions in the proof, but it is additionally supposed to return the i th position of the original string for given i .

For $\Sigma = \mathbb{F}_2^a$ for some $a \in \mathbb{N}$, let LIN_Σ denote the problem of solving linear equations where an instance consists of k variables that can have a value from Σ , and a system of linear equations C on $k \cdot a$ variables over \mathbb{F}_2 canonically represented by the k variables over Σ . It is equivalent to LIN over \mathbb{F}_2 on $k \cdot a$ variables, except that we consider each block of a variables as one variable that can take a value from Σ . We define a decoder for LIN_Σ below.

Definition 3.4 (Decoder for LIN_Σ). *Let $\Sigma = \mathbb{F}_2^a$ and $\sigma = \mathbb{F}_2^b$ for some a and b . A decoder for LIN_Σ over a proof alphabet σ with parameters $m, q, r : \mathbb{N} \rightarrow \mathbb{N}$ is a probabilistic polynomial-time algorithm \mathcal{D} . It is given a system of linear equations C on n variables over Σ , and an index $j \in [n]$*

as input, and oracle access to a proof π of length $m(n)$ over proof alphabet σ . It tosses $r = r(n)$ random coins and generates (1) a sequence of $q = q(n)$ locations $I = (i_1, \dots, i_q)$ and (2) a (local decoding) function $f : \sigma^q \rightarrow \Sigma \cup \{\perp\}$. \mathcal{D} is called linear if for every f , $P := f^{-1}(\Sigma)$ is an affine space of $\sigma^q = (\mathbb{F}_2^{qb})$ and $f : P \rightarrow \Sigma$ is an affine function over the base field \mathbb{F}_2 .

Now we define a dPCP for LIN_Σ . The dPCP in [DH13] is defined for `CIRCUIT SAT`, whereas ours is for LIN_Σ . Note that unlike in [DH13], the dPCP we will construct does not imply any computational hardness, because it only proves whether the given system of linear equations is perfectly satisfiable or not, which is a computationally easy problem. The key point is it proves the system is satisfiable using a proof which is in some sense “locally decodable”. The dPCP will then be composed with the previous linear robust PCP, which is a system of linear equations with *imperfect completeness*, to reduce the query complexity.

Definition 3.5 (Decodable PCPs for LIN_Σ). *For functions $\delta : \mathbb{N} \rightarrow [0, 1]$ and $L : \mathbb{N} \rightarrow \mathbb{N}$, we say that a PCP decoder \mathcal{D} is a decodable probabilistically checkable proof (dPCP) system for LIN_Σ with perfect completeness, soundness δ and list size L if the following completeness and soundness properties hold for every system of linear equations C on n variables over Σ .*

Completeness. *For any $y \in \Sigma^n$ that satisfies every equation in C , there exists a proof $\pi \in \sigma^m$, also called a decodable PCP, such that*

$$\Pr_{j,I,f} [f(\pi_I) = y_j] = 1,$$

where $j \in [n]$ is chosen uniformly at random and I, f are distributed according to C, j , and the verifier’s random coins.

Soundness. *For any $\pi \in \sigma^m$, there is a list of $0 \leq \ell \leq L$ strings y^1, \dots, y^ℓ , where each y^i satisfies all equations in C , such that*

$$\Pr_{j,I,f} [f(\pi_I) \notin \{\perp, y_j^1, \dots, y_j^\ell\}] \leq \delta.$$

Robust soundness. *We say that \mathcal{D} is a robust dPCP system for LIN_Σ with robust soundness error δ , if the soundness criterion above can be strengthened to the following robust soundness criterion,*

$$\mathbf{E}_{j,I,f} [\text{agr}(\pi_I, \text{BAD}(f))] \leq \delta,$$

where

$$\text{BAD}(f) := \{w \in \sigma^q : f(w) \notin \{\perp, y_j^1, \dots, y_j^\ell\}\}.$$

The dPCP result we use is the following.

Theorem 3.6 (dPCP, Analog of [DH13, Theorem 6.5]). *There exist constants $\alpha, \gamma > 0$ such that for every $\delta \geq n^{-\alpha}$ and input alphabet size Σ of size at most n^γ , LIN_Σ has a linear robust decodable PCP system with perfect completeness, robust soundness error $\delta > 0$ and list size $L \leq 2/\delta$, query complexity $n^{1/8}$, proof alphabet σ of size n^γ , proof length $\text{poly}(n)$, and randomness complexity $O(\log n)$.*

The proof of this theorem is identical to that of [DH13, Theorem 6.5], except that the initial starting point is LIN_Σ instead of $\text{CIRCUIT SAT}_\Sigma$. Since the starting point is linear and all transformations are linear, the final object is also linear. The perfect completeness is also maintained. As mentioned before, the dPCP constructed here does not imply any computational hardness unlike in [DH13].

3.3 Composition

After having building blocks, Dinur and Harsha [DH13] show how to compose those blocks iteratively to reduce the query complexity and the alphabet size. Each composition involves several other operations including alphabet and degree reductions. While the soundness analyses for them are already proved in [DH13], we show that all of their operations preserve linearity and robust completeness.

Efficient Composition ([DH13, Theorem 4.2]). In the composition, given a regular robust linear PCP verifier V and a robust linear PCP decoder \mathcal{D} , the composed verifier V' expects a decodable PCP for each constraint of V . Recall that the linearity of V is equivalent to the fact that each constraint of V is a system of linear equations over \mathbb{F}_2 , which is exactly what \mathcal{D} expects. An informal description of the composed verifier is as follows:

1. Randomly choose a location i of the proof for V . Let C_1, \dots, C_D be the constraints of V containing the location.
2. Using a $(\varepsilon, \varepsilon^2)$ -sampler $([D], [D], E)$ and a random $s \in [D]$, choose a subset $S \subseteq \{1, \dots, D\}$ and run the inner PCP decoder \mathcal{D} for each C_j with $j \in S$ to decode the i th symbol in the original proof.
3. Accept if all the values returned by the PCP decoders are the same.

For the second step above, we use $(\varepsilon, \varepsilon^2)$ -samplers given in [Gol11]. Theorem 4.2 of [DH13] shows the soundness of the composed verifier V' , yielding table 1 below (Table 4.2 in [DH13]).

	V	\mathcal{D}	V'
proof alphabet	Σ	σ	σ
randomness complexity	R	r	$\log M + r + \log D$
query complexity	Q	q	$4/\varepsilon^4 \cdot q$
proof degree	D	d	d
proof length	M	m	$2^R \cdot m$
robust soundness error	Δ	δ	$\Delta L + 4L\varepsilon + \delta$
list size	-	L	-

Table 1: Parameters for Composition.

We check this composition preserves robust completeness and linearity.

- **Linearity:** Linearity (over \mathbb{F}_2) is preserved if both V and \mathcal{D} are linear, since the only additional check we perform is to check whether the returned values are equal.

- **Robust completeness:** Suppose that there exists a proof Π for V that achieves the robust completeness of at least $1 - \xi$. Recall that the composed verifier expects, for each constraint of the outer PCP, a satisfying assignment encoded by the inner dPCP. The proof for the composed verifier is the concatenation of all these encodings. Consider the proof to the composed verifier constructed by the honest encoding of the assignment that achieves the robust completeness for the outer PCP verifier. We will show that this proof achieves robust completeness $1 - \xi$.

Let i be a proof location in the outer PCP and C_1, \dots, C_D be the constraints involving i . Furthermore, let ξ_i be the fraction of these constraints violated by the proof. Since Π is at least $(1 - \xi)$ -robustly complete, we have $\mathbf{E}_i[\xi] \leq \xi$. For each sample s chosen in the sampler, let $\xi_{i,s}$ be the fraction of constraints in S (chosen by sampler) that are violated. By regularity of sampler, we have $\mathbf{E}_s[\xi_{i,s}] \leq \xi_i$.

A local view of the composed verifier (corresponding to i, s and the inner dPCP randomness) comprises of the concatenation of the local views of the dPCP encodings corresponding to the constraints in S . Since the the inner dPCP has perfect completeness we have the following. Whenever the constraint is satisfied, the corresponding inner dPCP's encodings satisfies all constraints while we have no guarantee when the constraint is not satisfied. Since for each (i, s) , the fraction of violated constraints is $\xi_{i,s}$, we have that at least $(1 - \xi_{i,s})$ -fraction of the local inner views corresponding to (i, s) are satisfying and furthermore they all decode to the same $\Pi(i)$. Hence, the local view of the composed verifier corresponding to (i, s) is at least $(1 - \xi_{i,s})$ -close to a satisfying view. Hence, the robust completeness of this honest proof is at least $\mathbf{E}_{i,s}[1 - \xi_{i,s}] \geq 1 - \xi$.

3.4 Label Cover Operations

After the composition, the alphabet reduction step is applied to ensure that the alphabet size is polynomial in the query complexity and the inverse of the soundness. Also, since the basic robust PCP given in theorem 3.2 is not necessarily regular, we also need to show how to make the initial robust PCP regular. This subsection introduces various such operations and explains why they preserve robust completeness and linearity.

Degree Reduction ([DH13, Theorem 5.1]). Given an instance of LABEL COVER $G = (A, B, E)$, the degree reduction makes the instance right-regular by appropriately duplicating right vertices and each edge exactly the same number of times. Theorem 5.1 of [DH13] ensures that by increasing robust soundness by 4μ additively, we can ensure that the right degree is $4/\mu^4$ for all right vertices. We check that this operation preserves linearity and robust completeness.

- **Linearity:** Linearity is obviously preserved, because there is no change in the constraint.
- **Robust completeness:** Since each edge is duplicated the same number of times, robust completeness does not decrease.

Alphabet Reduction ([DH13, Theorem 5.5]). Given an instance of LABEL COVER $G = (A, B, E)$ where Σ_A and Σ_B are the alphabet set of the left (bigger) side and the right (smaller) side respectively, the alphabet reduction replaces Σ_B by a smaller set σ by finding a suitable linear

code $C : \Sigma_B \rightarrow \sigma^k$ and replacing each vertex $b \in B$ by k vertices b_1, \dots, b_k . Then assigning $x \in \Sigma_B$ to b corresponds to assigning $(C(x))_i$ to b_1, \dots, b_i . Theorem 5.5 of [DH13] ensures that if C has a relative distance $1 - \eta^3$, this operation increases robust soundness by at most 3η additively. We check that this operation preserves linearity and robust completeness.

- **Linearity:** Linearity over \mathbb{F}_2 is preserved if the code $C : \Sigma_B \rightarrow \sigma^k$ is linear with $\sigma = \mathbb{F}_{2^a}$ as the base field for some $a \in \mathbb{N}$. The code used in Remark 5.4 of [DH13] is already linear.
- **Robust completeness:** If an edge (a, b) of the original LABEL COVER instance is preserved and the new instance follows the honest encoding, all k edges of the new instance corresponding to (a, b) will be satisfied. Therefore, robust completeness cannot decrease.

Flip Sides ([DH13, Section 5.3]). Given an instance of LABEL COVER $G = (A, B, E)$ where each right vertex $b \in B$ has degree d , the flip side is achieved by flipping A and B , and assigning each $v \in B$ a label from Σ_A^d , which is supposed to denote the assignments to its neighbors in the original instance. If $v \in B$ has $u_1, \dots, u_d \in A$ as neighbors, (v, u_i) in the new instance is satisfied (i) if the label $(a_1, \dots, a_d) \in \Sigma_A^d$ for v has $b \in \Sigma_B$ such that the label pair (a_i, b) satisfies the edge (u_i, v) in the old instance, and (ii) if a_i is equal to the label assigned to u_i . This obviously does not change the robust soundness. We check that it also preserves linearity and robust completeness.

- **Linearity:** Linearity is preserved, because for each $v \in B$, the set of (a_1, \dots, a_d) satisfying (i) above is an affine subspace of $(\Sigma_A)^d$, and the new constraint is merely a projection.
- **Robust completeness:** Cannot decrease since if $v \in B$ was assigned $b \in \Sigma_B$ in the original instance, it can be assigned $(a_1, \dots, a_d) \in \Sigma_A^d$ such that (i) $\pi_{(u_i, v)}(a_i) = b$, and (ii) a_i was assigned to u_i if (u_i, v) was satisfied in the original instance.

We use a combination of the above 3 operations to get a regular LABEL COVER instance, as shown below.

Given an $\varepsilon > 0$, by using $(O(\varepsilon), O(\varepsilon^2))$ -samplers in the composition and doing the above operations with $\eta = O(\varepsilon)$, $d = O(1/\varepsilon^4)$, distance $1 - O(\varepsilon^3)$, $|\sigma| = O(1/\varepsilon^6)$, $k = O(1/\varepsilon^6) \cdot |\Sigma'| \leq O(1/\varepsilon^6) \cdot q|\Sigma|$, we can deduce the following lemma.

Lemma 3.7 ([DH13, Lemma 5.7]). *For all $\varepsilon : \mathbb{N} \rightarrow [0, 1]$, suppose L has a robust linear PCP verifier V with randomness complexity r , query complexity q , proof length m , average proof degree D_B , robust completeness c , robust soundness error δ over a proof alphabet Σ . Then L has a regular reduced linear robust PCP verifier, which we shall denote by $\text{regular}_\varepsilon(V)$ with*

- *randomness complexity $\log m + \log D_B$,*
- *query complexity $O(q \log |\Sigma| / \varepsilon^{14})$,*
- *proof length $O(q^2 2^r \log |\Sigma| / \varepsilon^{10})$,*
- *proof degree $O(1/\varepsilon^4)$,*
- *proof alphabet σ of size at most $O(1/\varepsilon^6)$,*
- *robust completeness c ,*
- *and robust soundness $\delta + \varepsilon$.*

LABEL COVER (Robust PCPs)	I	Degree Red. ($\rightarrow d$)	Flip	Degree Red. ($\rightarrow d$)	Alphabet Red. ($\rightarrow \sigma$)
# left vertices (randomness)	n	n	mD_B	mD_B	mD_B
# right vertices (proof length)	m	mD_B	n	$nD_A d$	$nD_A d k$
left degree (query complexity)	D_A^*	dD_A^*	d	d^2	$d^2 k$
right degree (proof degree)	D_B^*	d	$D_A d^*$	d	d
left alphabet (# accepting conf.)	Σ_A	Σ_A	Σ_A^d	Σ_A^d	Σ_A^d
right alphabet (proof alphabet)	Σ_B	Σ_B	Σ_A	Σ_A	σ
soundness error (rob. soundness error)	δ	$\delta + 4\mu$	$\delta + 4\mu$	$\delta + 8\mu$	$\delta + 8\mu + 3\eta$
rob. completeness (rob. completeness)	$1 - \xi$	$1 - \xi$	$1 - \xi$	$1 - \xi$	$1 - \xi$

Table 2: Sequence of steps to regularize the LABEL COVER instance. * denotes irregular instances where the number denotes the average degree.

3.5 Putting things together

Finally we prove theorem 1.8 on the hardness of LINEAR LABEL COVER. Let $c > 0$ be an arbitrary constant. Let \mathcal{D} be the PCP decoder from theorem 3.6 and \mathcal{V} be the robust PCP from theorem 3.2 with robust completeness $1 - \delta$ with $\delta = \log^c n$, robust soundness error $\varepsilon = 1/\log^{c_0} n$ for some $c_0 > 1$, query complexity $1/\varepsilon^{O(1)}$, randomness complexity $O(\log n)$ and proof length $\text{poly}(n)$.

Lemma 3.8 ([DH13, Lemma 6.6]). *Let \mathcal{D} , \mathcal{V} , ε, δ be as defined above and set $\varepsilon_i = (\varepsilon)^{1/3^i}$. There exist constants $c_0, c_1, c_3 > 0$ such that for every $i \geq 0$ as long as $\varepsilon_i < c_0$, the following holds. GAP LIN($1 - \delta, 0.9$) has a regular linear robust PCP verifier V_i with query complexity $1/\varepsilon_i^{c_1}$, robust completeness $1 - \delta$, robust soundness error $2\varepsilon_i$, proof alphabet Σ_i of size c_3/ε_i^6 , randomness complexity $O(\log n)$ and proof length $\text{poly}(n)$.*

Proof. The proof is similar to [DH13], and is a sequence of compositions. We start with the regularized robust verifier given by applying the sequence of steps given in section 3.4 to the robust PCP verifier given in theorem 3.2. In each subsequent step, we compose the robust verifier obtained in the previous step with a dPCP, and apply the alphabet reduction (Theorem 5.5 of [DH13]) to reduce the size of the alphabet to c_3/ε_{i+1}^6 . All the parameters remain the same as in [DH13], and we only need to focus on the two additional properties we need, linearity and robust completeness.

Recall that a PCP with robust completeness $1 - \delta$, when composed with a dPCP with perfect completeness, yields a composed PCP with robust completeness $1 - \delta$. In each step the inner PCP decoder has perfect completeness, therefore the robust completeness of the composed PCP is preserved. Recall that the alphabet reduction step also doesn't affect the perfect completeness.

Linearity is also preserved because all basic components are linear and all steps (e.g., composition, alphabet reduction, and regularization) preserve linearity as previously discussed. ■

The above lemma shows that we can iteratively reduce the query complexity until some absolute constant while maintaining the soundness and the alphabet size polynomial in the query complexity. (And the total size of the instance always remains polynomial in n .) Only a constant number of iterations is needed until $(\text{proof alphabet size})^{(\text{query complexity})}$, an upper bound on the size of alphabet in the equivalent LABEL COVER instance, becomes polynomial in n . This proves our main theorem 1.8 for LINEAR LABEL COVER.

Proof of theorem 1.8. Set i from lemma 3.8 so that

$$(\text{proof alphabet size})^{(\text{query complexity})} = (c_3/\varepsilon_i^6)^{1/\varepsilon_i^{c_1}} = \exp\left(\frac{1}{\varepsilon_i^{c_1}} \cdot \log\left(\frac{c_3}{\varepsilon_i^6}\right)\right) \leq \text{poly}(n).$$

This ensures that $\varepsilon_i = 1/\log^{c_4} n$ for some $c_4 > 0$. Using the equivalence between LABEL COVER and robust PCP, we have a hardness of LABEL COVER where the number of vertices and the size of label are bounded by $\text{poly}(n)$, and the completeness is at least $1 - 1/\log^c n$, the soundness is $1/\log^{c_4} n$. Applying the parallel repetition of [DS14] $O(c/c_4)$ times to reduce the soundness to $1/\log^c n$ finishes the proof. ■

4 Reduction from Linear Label Cover to 3LIN

In this section, we prove our main theorem 1.6 for 3-LIN. Recall that theorem 3.3 shows a randomized polynomial reduction from 3-SAT to GAP LINEAR LABEL COVER($1 - \log^c n, \log^c n$) for

any constant $c > 0$, where the number of vertices as well as the number of labels are bounded by a polynomial. Therefore, the following theorem finishes the proof of theorem 1.6. The main idea is to use Hadamard codes instead of long codes using the fact that the LABEL COVER instance is linear. A similar argument was used in [Kho01].

Lemma 4.1. *There is a polynomial time reduction from GAP LINEAR LABEL COVER($1 - \delta, s$) to GAP 3-LIN($1 - \delta, 1/2 + \sqrt{s}/2$), where the size of the 3-LIN instance is polynomial in the number of vertices and the size of label in the LABEL COVER instance.*

Proof. Let $G = (A, B, E), \Sigma_A, \Sigma_B, \{\pi_e\}_{e \in E}$ be an instance of GAP LINEAR LABEL COVER ($1 - \delta, s$). Moreover, since the label cover is linear, let the labels to left hand side vertices come from \mathbb{F}_2^ℓ and the right hand side vertices from \mathbb{F}_2^r , and the mapping on each edge is an affine mapping. Our reduction is described by the following test.

Test

- Consider an edge (u, v) . The labels $x \in \mathbb{F}_2^\ell, y \in \mathbb{F}_2^r$ corresponding to the vertices have to satisfy $x = Ay + b$.
- From the proof, we randomly sample the Hadamard code of x at location α , and that of y at locations β and $\beta + \gamma$, where $\gamma = A^T \cdot \alpha$.
- Check if $\langle \alpha, x \rangle + \langle \beta, y \rangle + \langle \beta + \gamma, y \rangle = \langle \alpha, b \rangle$

Completeness. In the completeness case, if the labels x, y satisfy the edge in the LINEAR LABEL COVER, then we can see that the test will pass.

$$\begin{aligned}
 & \langle \alpha, x \rangle + \langle \beta, y \rangle + \langle \beta + \gamma, y \rangle \\
 &= \langle \alpha, Ay \rangle + \langle \alpha, b \rangle + \langle \beta, y \rangle + \langle \beta + \gamma, y \rangle \\
 &= \langle \alpha, Ay \rangle + \langle \alpha, b \rangle + \langle A^T \alpha, y \rangle \\
 &= \langle \alpha, b \rangle
 \end{aligned}$$

Therefore, if $1 - \delta$ edges are satisfiable in the linear LABEL COVER, at least $1 - \delta$ fraction of 3LIN constraints are satisfied.

Soundness. Consider the case where at most s fraction of edges can be satisfied for any labeling in the LINEAR LABEL COVER. Let the Hadamard code encoding function for the left vertices be L and right vertices be R . Consider their Fourier transforms,

$$\begin{aligned}
 L(\alpha) &= \sum_x \hat{L}(x) \chi_x(\alpha) \\
 R(\beta) &= \sum_y \hat{R}(y) \chi_y(\beta)
 \end{aligned}$$

Let's fix an edge, and analyze the probability that the test will accept. We switch to a $-1,+1$ notation for convenience.

$$\begin{aligned}
\Pr[\text{Test accepts}] &= \Pr_{\alpha,\beta}[\langle\alpha, x\rangle + \langle\beta, y\rangle + \langle\beta + A^T\alpha, y\rangle + \langle\alpha, b\rangle = 0] \\
&= \Pr_{\alpha,\beta}[(-1)^{\langle\alpha, x\rangle + \langle\beta, y\rangle + \langle\beta + A^T\alpha, y\rangle + \langle\alpha, b\rangle} = 1] \\
&= \frac{1 + \mathbf{E}_{\alpha,\beta} [L(\alpha)R(\beta)R(\beta + A^T\alpha)(-1)^{\langle\alpha, b\rangle}]}{2}
\end{aligned}$$

Consider the expectation on the right hand side of the above equation.

$$\begin{aligned}
&\mathbf{E}_{\alpha,\beta} [L(\alpha)R(\beta)R(\beta + A^T\alpha)(-1)^{\langle\alpha, b\rangle}] \tag{3} \\
&\leq \sum_{x,y} \hat{L}(x)\hat{R}(y)^2 \mathbf{E}_{\alpha,\beta} [\chi_x(\alpha)\chi_y(\beta)\chi_z(\beta + A^T\alpha)(-1)^{\langle\alpha, b\rangle}] \\
&\leq \sum_{x,y,x=Ay+b} \hat{L}(x)\hat{R}(y)^2 \\
&\leq \sqrt{\sum_{x,y,x=Ay+b} \hat{R}(y)^2} \sqrt{\sum_{x,y,x=Ay+b} \hat{L}(x)^2\hat{R}(y)^2}
\end{aligned}$$

In the above equation, the first term is bounded by 1, and therefore,

$$(3) \leq \sqrt{\sum_{x,y,x=Ay+b} \hat{L}(x)^2\hat{R}(y)^2}$$

Consider a random assignment where a left vertex gets a label x with probability $\hat{L}(x)^2$ and a right vertex gets a label y with probability $\hat{R}(y)^2$. The probability that such a random assignment would satisfy the edge, and therefore the expected fraction of edges satisfied, is exactly

$$\sum_{x,y,x=Ay+b} \hat{L}(x)^2\hat{R}(y)^2$$

If at most s fraction of edges can be satisfied by any assignment, then

$$s \geq \sum_{x,y,x=Ay+b} \hat{L}(x)^2\hat{R}(y)^2 \geq (2 \cdot \Pr[\text{Test accepts}] - 1)^2$$

or

$$\Pr[\text{Test accepts}] \leq \frac{1}{2} + \frac{\sqrt{s}}{2}$$

Therefore, the expected fraction of 3LIN constraints satisfied is at most $\frac{1}{2} + \frac{\sqrt{s}}{2}$. ■

References

- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, May 1998. (Preliminary version in *33rd FOCS*, 1992).

- [AS98] Sanjeev Arora and Shmuel Safra. Probabilistic checking of proofs: A new characterization of NP. *J. ACM*, 45(1):70–122, January 1998. (Preliminary version in *33rd FOCS*, 1992).
- [DH13] Irit Dinur and Prahladh Harsha. Composition of low-error 2-query PCPs using decodable PCPs. *SIAM J. Comput.*, 42(6):2452–2486, 2013. (Preliminary version in *51st FOCS*, 2009).
- [DS14] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Proc. 46th ACM Symp. on Theory of Computing (STOC)*, pages 624–633, 2014.
- [Gol11] Oded Goldreich. A sample of samplers: A computational perspective on sampling. In Oded Goldreich, editor, *Studies in Complexity and Cryptography. Miscellanea on the Interplay between Randomness and Computation*, volume 6650 of *LNCS*, pages 302–332. Springer, 2011.
- [Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, July 2001. (Preliminary version in *29th STOC*, 1997).
- [HV04] Johan Håstad and Srinivasan Venkatesh. On the advantage over a random assignment. *Random Structures Algorithms*, 25(2):117–149, 2004. (Preliminary version in *34th STOC*, 2002).
- [Kho01] Subhash Khot. Improved inapproximability results for maxclique, chromatic number and approximate graph coloring. In *Proc. 42nd IEEE Symp. on Foundations of Comp. Science (FOCS)*, pages 600–609, 2001.
- [Kho10] Subhash Khot. Inapproximability results for computational problems on lattices. In Phong Q. Nguyen and Brigitte Vallée, editors, *The LLL Algorithm - Survey and Applications*, Information Security and Cryptography, pages 453–473. Springer, 2010.
- [KN08] Subhash Khot and Assaf Naor. Linear equations modulo 2 and the L_1 diameter of convex bodies. *SIAM J. Comput.*, 38(4):1448–1463, 2008. (Preliminary version in *48th FOCS*, 2007).
- [Mos15] Dana Moshkovitz. The Projection Games Conjecture and the NP-hardness of \ln -approximating Set-Cover. *Theory Comput.*, 11:221–235, 2015. (Preliminary version in *15th APPROX*, 2012).
- [MR10] Dana Moshkovitz and Ran Raz. Two-query PCP with subconstant error. *J. ACM*, 57(5), 2010. (Preliminary version in *49th FOCS*, 2008).
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, June 1998. (Preliminary version in *27th STOC*, 1995).