Vertex Cover Might be Hard to Approximate to within $2 - \varepsilon$

Subhash Khot * Oded Regev †

Abstract

Based on a conjecture regarding the power of unique 2-prover-1-round games presented in [Khot02], we show that vertex cover is hard to approximate within any constant factor better than 2. We actually show a stronger result, namely, based on the same conjecture, vertex cover on k-uniform hypergraphs is hard to approximate within any constant factor better than k.

1 Introduction

Minimum vertex cover is the problem of finding the smallest set of vertices that touches all the edges in a given graph. This is one of the most fundamental NP-complete problems. A simple 2-approximation algorithm exists for this problem: construct a maximal matching by greedily adding edges and then let the vertex cover contain both endpoints of each edge in the matching. It can be seen that the resulting set of vertices indeed touches all the edges and that its size is at most twice the size of the minimum vertex cover. However, despite considerable efforts, state of the art techniques can only achieve an approximation ratio of $2 - o(1)$ [16, 21].

Given this state of affairs, one might strongly suspect that vertex cover is NP-hard to approximate within $2-\varepsilon$ for any $\varepsilon > 0$. This is one of the major open questions in the field of approximation algorithms. In [18], Håstad showed that approximating vertex cover within constant factors less than $\frac{7}{6}$ is NP-hard. This factor was recently improved by Dinur and Safra [10] to 1.36. In a related result, Arora et al. [1] considered algorithms based on linear programming. They showed an integrality gap of $2 - \varepsilon$ for a large family of linear programs for vertex cover. This implies that many linear programming based algorithms cannot obtain an approximation ratio better than 2.

We also consider the more general problem of vertex cover on k-uniform hypergraphs. A k-uniform hypergraph $H = (V, E)$ consists of a set of vertices V and a collection E of k-element subsets of V called hyperedges (or simply edges). A *vertex cover* of H is a subset of vertices $S \subseteq V$ such that every hyperedge in E intersects S, i.e., $e \cap S \neq \emptyset$ for each $e \in E$. An *independent set* in H is a subset whose complement is a vertex cover, or in other words, a subset of vertices that contains no hyperedge entirely within it. The Ek-Vertex-Cover problem is the problem of finding a minimum size vertex cover in a k-uniform hypergraph. Notice that for $k = 2$, this problem is the same as the vertex cover problem on graphs. The simple algorithm presented before can be easily extended to k-uniform hypergraphs, achieving a factor k approximation. However, as before, the

*Department of Computer Science, Princeton University, Princeton, NJ 08544. E-Mail: `khot@cs.princeton.edu`. Supported by Prof. Sanjeev Arora's David and Lucile Packard Fellowship, NSF Grant CCR-0098180, and an NSF ITR grant

†Department of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. Supported by an Alon Fellowship, by the Binational Science Foundation, by the Israel Science Foundation, and by the EU Integrated Project QAP.

best approximation algorithms yield only a tiny improvement, achieving a $k - o(1)$ approximation ratio [16].

The first explicit hardness result shown for Ek-Vertex-Cover was due to Trevisan [29] who showed (among other results) an inapproximability factor of $k^{1/19}$. Holmerin [20] showed that E4-Vertex-Cover is NP-hard to approximate within $(2 - \varepsilon)$. Independently, Goldreich [13] showed a direct 'FGLSS'-type [11] reduction (involving no use of the long-code, a crucial component in most recent PCP constructions) attaining a hardness factor of $(2-\varepsilon)$ for Ek-Vertex-Cover for some constant k. More recently, Holmerin [19] showed that Ek-Vertex-Cover is NP-hard to approximate within $k^{1-\varepsilon}$, and also that it is NP-hard to approximate E3-Vertex-Cover within factor $(3/2 - \varepsilon)$. Dinur, Guruswami and Khot [7] gave a fairly simple proof of an $\Omega(k)$ hardness result for Ek-Vertex-Cover and a more complicated proof that shows a factor $(k - 3 - \varepsilon)$ hardness for Ek-Vertex-Cover. Finally, a recent paper by Dinur et al. [8] improves upon all previous results by showing a $(k-1-\varepsilon)$ hardness result.

With this recent progress on the Ek-Vertex-Cover problem, there is a strong reason to believe that it is NP-hard to approximate Ek-Vertex-Cover within $k - \varepsilon$ for every $k \geq 2$. The current techniques, however, seem very inadequate to prove such a result. In [22], Khot presented the *unique games conjecture* as an approach to attack many fundamental open problems. The conjecture deals with 2-prover-1-round games where two (all-powerful) provers try to convince a probabilistic verifier that a certain NP-statement is true. The proof system is 1-round, meaning the verifier asks both the provers one question each and accepts or rejects depending on the provers' answers. The game is called *unique* if the answer of one prover completely determines the answer of the second prover and vice versa. The conjecture essentially states that it is NP-hard to distinguish whether the success probability of the provers' optimal strategy in a unique 2-prover-1-round game is very close to 1 or very close to 0. Assuming this conjecture, Khot was able to show several new hardness results including the hardness of the Min-2SAT-Deletion problem. He also observed that a variant of his conjecture would imply a $\sqrt{2} - \varepsilon$ hardness result for vertex cover.

In this paper, we continue this line of research and, assuming the unique games conjecture, we prove a tight $k - \varepsilon$ hardness result for Ek-Vertex-Cover. We obtain this by showing that given a k-uniform hypergraph that has an independent set of size $1-\frac{1}{k}-\varepsilon$, it is hard to find an independent set of size ε. We remark that for the case $k = 2$ we obtain a $2 - \varepsilon$ hardness result for vertex cover, giving further evidence that the factor 2 may be the right answer for this problem.

Main techniques: Many of the recent hardness results are shown via constructions of new Probabilistically Checkable Proof systems (PCPs) (see, e.g., [4, 18, 17, 15]). These constructions typically involve two modules, the so-called *Outer PCP* and the so-called *Inner PCP*. The Outer PCP is essentially a 2-prover-1-round game and the Inner PCP is based on long codes and often, the Fourier analysis of long codes. In almost all of the constructions, the Outer PCP is obtained from the PCP theorem [3, 2] together with Raz's parallel repetition theorem [27].

However, this standard recipe hasn't been very successful in attacking the vertex cover problem. Håstad's $\frac{7}{6}$ hardness remained the best known result for a long time. Dinur and Safra [10] were able to break this barrier by relying on techniques from extremal combinatorics. However, their approach still doesn't succeed in getting a hardness factor better than 1.36. Khot [22] observed that the bottleneck in getting hardness results for vertex cover and a number of other problems might be in the Outer PCP, a component which has remained untouched so far. His conjecture basically states that a strong enough Outer PCP exists. On top of such strong Outer PCPs, one can build Inner PCPs that yield the desired hardness of approximation results.

Khot's conjecture looks quite promising, at least in light of the lack of any other techniques. We think that it is worthwhile to investigate which problems could be solved via this conjecture and we show that vertex cover is one such problem. We combine this conjecture with the techniques from Dinur and Safra's paper [10], which include the biased long code, Friedgut's Theorem, and theorems in extremal set theory.

It turns out that we need the unique games conjecture in a stronger form than what is stated in [22]. A significant contribution of this paper is to show that the stronger form follows from the original form. Roughly speaking, the original form states that in the good case the provers in the 2-prover-1-round game have a strategy that convinces the verifier with probability close to 1. In the stronger form, the provers have a strategy such that the verifier accepts whenever both questions fall inside some set that contains almost all possible questions. A more precise description will be given later.

Discussion: Recently, many new hardness results have been proved assuming the unique games conjecture: Khot et al. [23] prove an optimal hardness result of roughly 0.878 for MAX-CUT; Chawla et al. [6] and independently, Khot and Vishnoi [24] prove super-constant hardness results for the Sparsest Cut and Multi-Cut problems; Dinur, Mossel, and Regev [9] prove that a variant of the unique games conjecture implies that it is NP-hard to color 3-colorable graphs with any constant number of colors. Most of these results are based on our strong form of the unique games conjecture or variants of it. The fact that the unique games conjecture implies so many hardness results in a unifying way can be taken as an evidence towards its truth. Further evidence is given by Khot and Vishnoi [24] who prove an $(1-o(1), o(1))$ integrality gap for a semidefinite programming relaxation of the problem underlying the unique games conjecture.

One possible way to disprove the unique games conjecture would be to find a polynomial time algorithm for the problem underlying the conjecture. Several such algorithms have been suggested recently, see Trevisan [30], Gupta and Talwar [14], and Charikar et al. [5]. However, none of these algorithms is strong enough to disprove the conjecture.

Finally, we mention that work on the unique games conjecture has led to unconditional results in Fourier analysis (the Majority is Stablest Theorem [26]) and lower bounds in metric embeddings (the disproval of a conjecture of Goemans and Linial that negative type metrics embed into ℓ_1 with constant distortion, see [24]). These results, being independent of the conjecture, indicate that research on the conjecture is worthwhile, even if the conjecture eventually turns out to be false.

It remains an important open problem to resolve the unique games conjecture. It would also be interesting to see further implications of it towards hardness of approximation results.

Overview of the paper: In Section 2, we describe the unique games conjecture, introduce tools for the analysis of set-families, and some theorems from extremal combinatorics. Section 3 explains the reduction to the stronger form of the conjecture and it is the crux of the paper. Section 4 explains the reduction to hypergraph vertex cover and shares many ideas with previous work such as [10] and [8].

2 Preliminaries

2.1 The label cover problem

For convenience, from now on we will adopt a more combinatorial terminology, and describe 2-prover-1-round games as instances of the *label cover problem*. We represent the provers' strategy by a labeling to a set of variables, one variable for each possible question. The verifier is represented by a probability distribution on pairs of variables along with a relation for each pair, specifying the acceptance criterion. More formally, an instance of the (bipartite, weighted) label cover problem is specified by a tuple $\Phi = (X, Y, R, \Psi, W)$. The sets X and Y contain variables, and we often refer to variables in X as *left* variables and to variables in Y as *right* vertices. The set R is the set of possible labels. For each $x \in X$, $y \in Y$, Ψ contains one relation $\psi_{xy} \subseteq R \times R$ and W contains its weight $w_{xy} \geq 0$. A *labeling* is a function L mapping $X \cup Y$ to R. A constraint ψ_{xy} is said to be *satisfied* by a labeling L if $(L(x), L(y)) \in \psi_{xy}$.

We denote by $w(\Phi, x)$ the sum $\sum_{y \in Y} w_{xy}$ and by $w(\Phi)$ the sum $\sum_{x \in X, y \in Y} w_{xy}$. Also, for a labeling L, the weight of satisfied constraints, denoted by $w_L(\Phi)$, is the sum of w_{xy} over all $x \in X$ and $y \in Y$ such that ψ_{xy} is satisfied by L. Similarly, we define $w_L(\Phi, x)$ as the sum of w_{xy} over all $y \in Y$ such that ψ_{xy} is satisfied by L.

The PCP theorem of [3, 2], together with the parallel repetition theorem of [27], show that the label cover problem is NP-hard in the following strong sense.

Theorem 2.1 ([3, 2, 27]) *For any $\gamma > 0$ there exists a $|R|$ such that the following is NP-hard. Given a bipartite weighted label cover instance Φ with label set R and with $w(\Phi) = 1$, distinguish between the following two cases:*

- *(YES Case): There exists a labeling L such that $w_L(\Phi) = 1$.*
- *(NO Case): For any labeling L, $w_L(\Phi) \leq \gamma$.*

In other words, it is hard to distinguish between the case where there exists a labeling that satisfies all constraints, and the case where no labeling satisfies more than a tiny fraction of constraints. This theorem is at the core of many recent NP-hardness results, including [18, 10].

However, as mentioned before, for the vertex cover problem (as well as several other problems), constructions based on Theorem 2.1 have failed to yield satisfactory results. To this end, Khot [22] introduced the unique games conjecture. Essentially, it says that even if we require all constraints in Ψ to have a very specific form, the problem is still NP-hard. More precisely, we say that a constraint $\psi_{xy} \in \Psi$ is *unique* if for each $a \in R$ there exists a unique $b \in R$ such that $(a, b) \in \psi_{xy}$ and vice versa; in other words, ψ_{xy} can be thought of as a matching between labels of x and labels of y. We say that the instance Φ is unique of all its constraints are unique. The unique games conjecture of [22] is the following.

Conjecture 2.2 (Bipartite Weighted Unique Games Conjecture) *For any $\zeta, \gamma > 0$ there exists a $|R|$ such that the following is NP-hard. Given a bipartite weighted unique label cover instance Φ with label set R and with $w(\Phi) = 1$, distinguish between the following two cases:*

- *(YES Case): There exists a labeling L such that $w_L(\Phi) \geq 1 - \zeta$.*
- *(NO Case): For any labeling L, $w_L(\Phi) \leq \gamma$.*

Note that we assume $\zeta > 0$ for otherwise the problem can be seen to be solvable in polynomial time.

2.2 On set families

For a set R, let $P(R)$ denote its power set, i.e., the family of all subsets of R. For a "bias parameter" $0 < p < 1$, we define the weight $\mu_p^R(F)$ of a set F as

$$\mu_p^R(F) \stackrel{def}{=} p^{|F|}(1-p)^{|R \setminus F|}.$$

We omit the superscript R when no confusion is possible. The weight of a family $\mathcal{F} \subseteq P(R)$ is defined as

$$\mu_p^R(\mathcal{F}) \stackrel{def}{=} \sum_{F \in \mathcal{F}} \mu_p^R(F).$$

Note that μ_p^R is a probability measure on $P(R)$. In order to choose a set F from the corresponding distribution, independently include in F each element of R with probability p. Hence a 'typical' set chosen from this distribution is of size roughly $p|R|$.

For a family $\mathcal{F}$, an element $\sigma \in R$ and a bias parameter p we define the *influence of* σ *on* $\mathcal{F}$ as

$$\mathrm{Inf}_p^R(\mathcal{F}, \sigma) \stackrel{def}{=} \mu_p^R(\{F \subseteq R \mid \text{exactly one of } F \cup \{\sigma\},\ F \setminus \{\sigma\} \text{ is in } \mathcal{F}\}).$$

As before, the superscript will often be omitted. In words, the influence of σ on $\mathcal{F}$ is the probability that for a random F chosen according to μ_p^R, σ 'affects' the containment of F in $\mathcal{F}$ (in the sense that exactly one of $F \cup \{\sigma\},\ F \setminus \{\sigma\}$ is in $\mathcal{F}$). The *average sensitivity* of a family is defined as the sum of the influences of all elements,

$$\mathrm{as}_p(\mathcal{F}) \stackrel{def}{=} \sum_{\sigma \in R} \mathrm{Inf}_p(\mathcal{F}, \sigma).$$

2.2.1 Monotone families and the Russo-Margulis theorem

A family $\mathcal{F} \subseteq P(R)$ is called *monotone* if for any $F' \subseteq F \subseteq R$, $F' \in \mathcal{F}$ implies $F \in \mathcal{F}$. Also, for any family $\mathcal{F} \subseteq P(R)$ we define its *monotone extension* as the family $\{F \subseteq R \mid \exists F' \subseteq F \text{ s.t. } F' \in \mathcal{F}\}$. It is easy to see that the latter is a monotone family that contains $\mathcal{F}$.

For a monotone family $\mathcal{F}$, one would expect $\mu_p^R(\mathcal{F})$ to be a non-decreasing function of p. Indeed, this follows from the following theorem, which also shows that the derivative of $\mu_p^R(\mathcal{F})$ is given by the average sensitivity.

Theorem 2.3 (Russo-Margulis Theorem [28, 25]) *If $\mathcal{F} \subseteq P(R)$ is a monotone family, then $\mu_p(\mathcal{F})$ is a non-decreasing and differentiable function of p and*

$$\frac{d\mu_p(\mathcal{F})}{dp} = \mathrm{as}_p(\mathcal{F}).$$

2.2.2 Friedgut's theorem

Definition 2.4 *A family $\mathcal{F} \subseteq P(R)$ is called a core-family with a core $C \subseteq R$ if there exists a family $\mathcal{H} \subseteq P(C)$ such that*

$$\forall\ F \in P(R), \quad F \in \mathcal{F} \ \textit{ iff } \ F \cap C \in \mathcal{H}.$$

In other words, $\mathcal{F}$ is a core family with core C if and only if the containment of a set F in $\mathcal{F}$ depends only on $F \cap C$. An important theorem of Friedgut states that every family with low average sensitivity is well-approximated by a core family with small core, where by 'small' we mean that its size does not depend on the size of the universe R.[1]

[1] In [12], this theorem is stated in the equivalent formulation of Boolean functions that depend on a few coordinates. The set of these coordinates is precisely the core.

Theorem 2.5 (Friedgut's Theorem [12]) *Let $\mathcal{F} \subseteq P(R)$ be a family and p be a bias parameter. Let $k = \mathrm{as}_p(\mathcal{F})$ and $\eta > 0$ be an accuracy parameter. Then there exists a core family $\widehat{\mathcal{F}}$ with core of size at most $c^{k/\eta}$ and $\mu_p(\mathcal{F}\Delta\widehat{\mathcal{F}}) \leq \eta$. Here $c > 0$ is a constant that depends only on p and Δ denotes the symmetric difference of the two families.*

It turns out that for monotone families, the requirement of low average sensitivity can be avoided if we are willing to slightly shift the bias parameter p. The precise statement of this result appears in the following theorem. It follows by combining the Russo-Margulis Theorem and Friedgut's Theorem, as was done in [10]. We include the proof for completeness.

Theorem 2.6 ([10]) *Let $0 < p < 1$ and $\varepsilon, \eta > 0$ be some reals. Then, for any monotone family $\mathcal{F} \subseteq P(R)$ there exists $p' \in (p, p+\varepsilon)$ and a core family $\widehat{\mathcal{F}} \subseteq P(R)$ with a core $C \subseteq R$ such that*

- *The average sensitivity of the family $\mathcal{F}$ w.r.t. the bias p' is at most $\frac{1}{\varepsilon}$, i.e., $\mathrm{as}_{p'}(\mathcal{F}) \leq \frac{1}{\varepsilon}$.*
- *The size of C depends only on p, ε, η.*
- *$\mu_{p'}(\mathcal{F}\Delta\widehat{\mathcal{F}}) < \eta$ where Δ denotes the symmetric difference of the two families.*

Proof: By Theorem 2.3 we have

$$\frac{d\mu_q(\mathcal{F})}{dq} = \mathrm{as}_q(\mathcal{F}).$$

Therefore, by the Mean-Value Theorem, there exists $p' \in (p, p+\varepsilon)$ such that

$$\mathrm{as}_{p'}(\mathcal{F}) = \left.\frac{d\mu_q(\mathcal{F})}{dq}\right|_{q=p'} = \frac{\mu_{p+\varepsilon}(\mathcal{F}) - \mu_p(\mathcal{F})}{\varepsilon} \leq \frac{1}{\varepsilon}.$$

Applying Friedgut's Theorem, we conclude that $\mathcal{F}$ is well-approximated by a core family of small core with respect to the bias parameter p'. ■

2.2.3 Two lemmas

We will need the following two lemmas. The first is similar to a lemma in [10]. The second follows from a theorem of Frankl and can be found as Lemma A.4 in [7]. For completeness, we include both proofs here.

Lemma 2.7 *Let $\mathcal{F} \subseteq P(R)$ be a monotone family and let $\eta > 0$, $0 < p < 1$ be some reals. Let $T \subseteq R$ be such that for every element $\sigma \in T$, $\mathrm{Inf}_p(\mathcal{F}, \sigma) < \eta$. Define a subfamily $\mathcal{F}'$ of the family $\mathcal{F}$ as*

$$\mathcal{F}' \stackrel{def}{=} \{F \in \mathcal{F} \mid F \setminus T \in \mathcal{F}\}.$$

Then,

$$\mu_p^R(\mathcal{F}') \geq \mu_p^R(\mathcal{F}) - \eta\ |T|\ (\min(p, 1-p))^{-|T|}.$$

Proof: Consider the family

$$\mathcal{F}'' \stackrel{def}{=} \{F \subseteq R \setminus T \mid F \cup T \in \mathcal{F},\ F \notin \mathcal{F}\}.$$

By examining the definition of μ_p^R, it can be seen that

$$\mu_p^R(\mathcal{F}) - \mu_p^R(\mathcal{F}') \leq \mu_p^{R \setminus T}(\mathcal{F}'').$$

By the definition of $\mathcal{F}''$, we have that for any set $F \in \mathcal{F}''$ there exists some $D \subseteq T$ and an element $\sigma \in T$ such that $F \cup D \cup \{\sigma\} \in \mathcal{F}$ but $F \cup D \notin \mathcal{F}$. Hence, any set $F \in \mathcal{F}''$ contributes at least $\mu_p^{R \setminus T}(F) \cdot \min(p, 1-p)^{|T|}$ to $\mathrm{Inf}_p^R(\mathcal{F}, \sigma)$ for some $\sigma \in T$. It remains to notice that the total influence of elements in T is at most $|T| \cdot \eta$. ■

For our second lemma, we will need the following theorem of Frankl.

Theorem 2.8 *Let $\mathcal{F} \subseteq P(R)$ where $|R| = n$ and every set in the family $\mathcal{F}$ has size m. Assume that every k sets in the family have nonempty intersection and $n > mk/(k-1)$. Then*

$$|\mathcal{F}| \leq \binom{n-1}{m-1}$$

Note that the family of all sets of size m containing one fixed element has size $\binom{n-1}{m-1}$.

Lemma 2.9 *Let $\varepsilon > 0$ be an arbitrarily small constant, $k \geq 2$ some integer, and define $p = 1 - \frac{1}{k} - \varepsilon$. Then, for a sufficiently large universe R, the following holds. For any $\mathcal{F} \subseteq P(R)$ such that $\mu_p(\mathcal{F}) \geq 1 - \frac{1}{k}$ there exist k sets in the family $\mathcal{F}$ whose intersection is empty.*

Proof: Let $n = |R|$ be the size of the universe. Assume on the contrary that every k sets in the family $\mathcal{F}$ have nonempty intersection. Partition the family $\mathcal{F}$ according to different set-sizes.

$$\mathcal{F}_i \stackrel{def}{=} \{F \mid F \in \mathcal{F},\ |F| = i\}.$$

With the bias parameter p, the total weight of all sets of size at least $(p+\varepsilon)n$ is less than ε when the universe is large enough. Hence

$$\mu_p(\mathcal{F}) \ < \ \varepsilon + \sum_{m < (p+\varepsilon)n} \mu_p(\mathcal{F}_m).$$

For $m < (p+\varepsilon)n$, we have $n > mk/(k-1)$. Since every k sets in the family $\mathcal{F}_m$ have a nonempty intersection, applying Frankl's Theorem, we get

$$|\mathcal{F}_m| \ \leq \ \binom{n-1}{m-1}.$$

Noting that every set in $\mathcal{F}_m$ has weight $p^m(1-p)^{n-m}$ we have

$$\begin{aligned}
\mu_p(\mathcal{F}) \ &< \ \varepsilon + \sum_{m<(p+\varepsilon)n} \binom{n-1}{m-1} p^m (1-p)^{n-m} \\
&\leq \ \varepsilon + p \left(\sum_m \binom{n-1}{m-1} p^{m-1} (1-p)^{(n-1)-(m-1)} \right) \\
&= \ \varepsilon + p = 1 - \frac{1}{k}
\end{aligned}$$

which gives a contradiction. ■

3 Strong Unique Games Conjecture

Our goal in this section is to describe the strong form of the unique games conjecture, and prove that it follows from the original conjecture. But first, we describe two variants of the label cover problem. The first variant we consider is that of *unweighted* label cover instances. Such an instance is given by a tuple $\Phi = (X, Y, R, \Psi, E)$. The multiset E includes pairs $(x, y) \in X \times Y$ and we can think of (X, Y, E) as a bipartite graph (possibly with parallel edges). For each $(x, y) \in E$, Ψ includes a constraint ψ_{xy}.

In our second variant, the instances are (possibly) not bipartite. Such instances are given by a tuple $\Phi = (X, R, \Psi, E)$ where X is a set of vertices, E is a multiset of pairs $(x_1, x_2) \in X \times X$, and Ψ includes a constraint $\psi_{x_1x_2}$ for each $(x_1, x_2) \in E$. Here, a labeling is a function from X to R.

Finally, we define a t-labeling (in the non-bipartite case) as a function L that labels each variable $x \in X$ with a set of values $L(x) \subseteq R$ such that $|L(x)| \leq t$ for all $x \in X$. A t-labeling L is said to *satisfy* a constraint $\psi_{x_1x_2}$ if and only if there exists $a \in L(x_1)$, $b \in L(x_2)$ such that $(a, b) \in \psi_{x_1x_2}$.

Conjecture 3.1 (Strong Unique Games Conjecture) *For any $\zeta, \gamma > 0$ and $t \in \mathbb{N}$ there exists some $|R|$ such that the following is NP-hard. Given a non-bipartite unweighted unique label cover instance $\Phi = (X, R, \Psi, E)$ distinguish between the following two cases:*

- *(YES Case): There exists a labeling L and a set $X_0 \subseteq X$, $|X_0| \geq (1-\zeta)|X|$, such that L satisfies all constraints between variables of X_0.*
- *(NO Case): For any t-labeling L and any set $X_0 \subseteq X$, $|X_0| \geq \gamma|X|$, not all constraints between variables of X_0 are satisfied by L.*

Our aim in this section is to prove the following theorem.

Theorem 3.2 *Conjecture 2.2 implies Conjecture 3.1.*

The proof follows by combining Lemmas 3.3, 3.4, 3.6, and 3.8. Each lemma presents an elementary transformation between variants of the label cover problem. The first transformation creates a (weighted, bipartite, unique) label cover instance in which all the X variables have the same weight.

Lemma 3.3 *There exists an efficient procedure that given a weighted bipartite unique label cover instance $\Phi = (X, Y, R, \Psi, W)$ with $w(\Phi) = 1$ and a constant ℓ, outputs a weighted bipartite unique label cover instance $\Phi' = (X', Y, R, \Psi', W')$ with the following properties:*

- *For all $x' \in X'$, $w(\Phi', x') = 1$.*
- *For any $\zeta \geq 0$, if there exists a labeling L to Φ such that $w_L(\Phi) \geq 1 - \zeta$ then there exists a labeling L' to Φ' in which $1 - \sqrt{(1 + \frac{1}{\ell-1})\zeta}$ of the variables x' in X' satisfy that $w_{L'}(\Phi', x') \geq 1 - \sqrt{(1 + \frac{1}{\ell-1})\zeta}$.*
- *For any $\beta, \gamma > 0$, if there exists a labeling L' to Φ' in which β of the variables x' in X' satisfy $w_{L'}(\Phi', x') \geq \gamma$, then there exists a labeling L to Φ such that $w_L(\Phi) \geq (1 - \frac{1}{\ell})\beta\gamma$.*

Proof: Given Φ as above, we define $\Phi' = (X', Y, R, \Psi', W')$ as follows. The set X' includes $k(x)$ copies of each $x \in X$, $x^{(1)}, \ldots, x^{(k(x))}$ where $k(x)$ is defined as $\lfloor \ell \cdot |X| \cdot w(\Phi, x) \rfloor$. For every $x \in X$, $y \in Y$ and $i \in \{1, \ldots, k(x)\}$ we define $\psi'_{x^{(i)}y}$ as ψ_{xy} and the weight $w'_{x^{(i)}y}$ as $w_{xy}/w(\Phi, x)$. Notice that $w(\Phi', x') = 1$ for all $x' \in X'$ and that $(\ell - 1)|X| \le |X'| \le \ell|X|$ since $w(\Phi) = 1$. Moreover, for any $x \in X$, $y \in Y$, the total weight of constraints created from ψ_{xy} is $k(x)w_{xy}/w(\Phi, x) \le \ell|X|w_{xy}$.

We now prove the second property. Given a labeling L to Φ that satisfies constraints of weight at least $1-\zeta$, consider the labeling L' defined by $L'(x^{(i)}) = L(x)$ and $L'(y) = L(y)$. By the property mentioned above, the total weight of unsatisfied constraints in Φ' is at most $\ell|X|\zeta$. Since the total weight in Φ' is at least $(\ell - 1)|X|$, we obtain that the fraction of unsatisfied constraints is at most $(1 + \frac{1}{\ell-1})\zeta$. Hence, by a Markov argument, we obtain that for at least $1 - \sqrt{(1 + \frac{1}{\ell-1})\zeta}$ of the X' variables $w_{L'}(\Phi', x) \ge 1 - \sqrt{(1 + \frac{1}{\ell-1})\zeta}$.

We now prove the third property. Assume we are given a labeling L' to Φ' for which β of the variables satisfy $w_{L'}(\Phi', x') \ge \gamma$. We claim that this implies that there exists a labeling L'' to Φ' for which β of the variables satisfy $w_{L''}(\Phi', x') \ge \gamma$ and moreover, for every $x \in X$, $L''(x^{(1)}), \ldots, L''(x^{(k(x))})$ are all the same. Indeed, this holds since the constraints between $x^{(i)}$ and the Y variables are the same for all $i \in \{1, \ldots, k(x)\}$ so we can define $L''(x^{(i)})$ as the 'best' labeling among $L'(x^{(1)}), \ldots, L'(x^{(k(x))})$. We now define the labeling L as $L(x) = L''(x^{(1)})$. The weight of constraints satisfied by L is

$$\begin{aligned}\sum_{x \in X} w_L(\Phi, x) &\ge \frac{1}{\ell|X|} \sum_{x \in X} k(x) \cdot w_L(\Phi, x)/w(\Phi, x) \\ &= \frac{1}{\ell|X|} \sum_{x' \in X'} w_{L''}(\Phi', x') \\ &\ge \frac{1}{\ell|X|} \beta|X'|\gamma \ge \left(1 - \frac{1}{\ell}\right)\beta\gamma\end{aligned}$$

where the first inequality follows from the definition of $k(x)$. ■

The second transformation creates an *unweighted* label cover instance. Moreover, the instances created by this transformation are left-regular, in the sense that the number of constraints $(x, y) \in E$ incident to each $x \in X$ is the same.

Lemma 3.4 *There exists an efficient procedure that given a constant ℓ and a weighted bipartite unique label cover instance $\Phi = (X, Y, R, \Psi, W)$ with $w(\Phi, x) = 1$ for all $x \in X$, outputs an* unweighted *bipartite unique label cover instance $\Phi' = (X, Y, R, \Psi', E')$ with the following properties:*

- *All left degrees are equal to $\alpha = \ell|Y|$.*
- *For any $\beta, \zeta > 0$, if there exists a labeling L to Φ such that $w_L(\Phi, x) \ge 1 - \zeta$ for at least $1 - \beta$ of the variables in X, then there exists a labeling L' to Φ' in which for at least $1 - \beta$ of the variables in X, at least $1 - \zeta - 1/\ell$ of their incident constraints are satisfied.*
- *For any $\beta, \gamma > 0$, if there exists a labeling L' to Φ' in which β of the variables in X have at least γ of their incident constraints satisfied, then there exists a labeling L to Φ such that for β of the variables in X, $w_L(\Phi, x) > \gamma - 1/\ell$.*

Proof: We define the instance $\Phi' = (X, Y, R, \Psi', E')$ as follows. For each $x \in X$, choose some $y_0(x) \in Y$ such that $w_{xy_0(x)} > 0$. For every $x \in X$, $y \ne y_0(x)$, E' contains $\lfloor \alpha w_{xy} \rfloor$ edges from x to y associated with the constraint ψ_{xy}. Moreover, for every $x \in X$, E' contains $\alpha - \sum_{y \in Y \setminus \{y_0(x)\}} \lfloor \alpha w_{xy} \rfloor$ edges from x to $y_0(x)$ associated with the constraint $\psi_{xy_0(x)}$. Notice that all left degrees are equal to α. Moreover, for any x, $y \ne y_0(x)$, we have that the number of edges between x and y is at most αw_{xy} and the number of edges from x to $y_0(x)$ is at most $\alpha w_{xy_0(x)} + |Y| = \alpha(w_{xy_0(x)} + 1/\ell)$.

Consider a labeling L to Φ and let $x \in X$ be such that $w_L(\Phi, x) \geq 1 - \zeta$. Then, in Φ', the same labeling satisfies that the number of incident constraints to x that are satisfied is at least $(1-\zeta-1/\ell)\alpha$. Finally, consider a labeling L' to Φ' and let $x \in X$ have γ of its incident constraints satisfied. Then, $w_{L'}(\Phi, x) > \gamma - \frac{1}{\ell}$. ■

In the third lemma we modify a left-regular unweighted label cover instance so that it has the following property: if there exists a labeling to the original instance that for many variables satisfies many of their incident constraints, then the resulting instance has a labeling that for many variables satisfies *all* their incident constraints. But first, we prove a combinatorial claim.

Claim 3.5 *For any integer ℓ, finite set R, and real $0 < \gamma < \frac{1}{\ell^2}$, let $F \subseteq R$ be a multiset with the property that no element $i \in R$ appears more than $\gamma|F|$ times in F. Then, the probability that a sequence of elements $i_1, i_2, \ldots, i_\ell$ chosen uniformly from F (with repetitions) contains no two identical elements is at least $1 - \ell^2\gamma$.*

Proof: By the union bound, it suffices to prove that $\Pr[i_1 = i_2] \leq \gamma$. This follows by fixing i_1 and using the assumption on F. ■

Lemma 3.6 *There exists an efficient procedure that given an unweighted bipartite unique label cover instance $\Phi = (X, Y, R, \Psi, E)$ with all left-degrees equal to some α, and a constant ℓ, outputs an unweighted bipartite unique label cover instance $\Phi' = (X', Y, R, \Psi', E')$ with the following properties:*

- *All left degrees are equal to ℓ.*
- *For any $\beta, \zeta \geq 0$, if there exists a labeling L to Φ such that for at least $1 - \beta$ of the variables in X $1 - \zeta$ of their incident constraints are satisfied, then there exists a labeling L' to Φ' in which $(1-\zeta)^\ell(1-\beta)$ of the X' variables have all their ℓ constraints satisfied.*
- *For any $\beta > 0$, $0 < \gamma < \frac{1}{\ell^2}$, if in any labeling L to Φ at most β of the variables have γ of their incident constraints satisfied, then in any labeling L' to Φ', the fraction of satisfied constraints is at most $\beta + \frac{1}{\ell} + (1-\beta)\ell^2\gamma$.*

Proof: We define $\Phi' = (X', Y, R, \Psi', E')$ as follows. For each $x \in X$, consider its neighbors $(y_1, \ldots, y_\alpha)$ listed with multiplicities. For each sequence $(y_{i_1}, \ldots, y_{i_\ell})$ where $i_1, \ldots, i_\ell \in \{1, \ldots, \alpha\}$ we create a variable in X'. This variable is connected to $y_{i_1}, \ldots, y_{i_\ell}$ with the same constraints as x, namely $\psi_{xy_{i_1}}, \ldots, \psi_{xy_{i_\ell}}$. Notice that the total number of variables created from each $x \in X$ is α^ℓ. Hence, $|X'| = \alpha^\ell|X|$.

We now prove the second property. Assume that L is a labeling to Φ such that for at least $1 - \beta$ of the variables in X, $1 - \zeta$ of their incident constraints are satisfied. Let L' be the labeling to Φ' assigning to each of the variables created from $x \in X$ the value $L(x)$ and for each $y \in Y$ the value $L(y)$. Consider a variable $x \in X$ that has $1 - \zeta$ of its incident constraints satisfied and let Y_x denote the set of variables $y \in Y$ such that ψ_{xy} is satisfied. Then among the variables in X' created from x, the number of variables that are connected only to variables in Y_x is at least $\alpha^\ell(1-\zeta)^\ell$. Therefore, the total number of variables all of whose constraints are satisfied by L' is at least

$$\alpha^\ell(1-\zeta)^\ell(1-\beta)|X| = (1-\zeta)^\ell(1-\beta)|X'|.$$

We now prove the third property. Assume that in any labeling L to Φ at most β of the X variables have γ of their incident constraints satisfied. Let L' be an arbitrary labeling to Φ'. For each $x \in X$ define $F_x \subseteq R$ as the multiset that contains for each constraint incident to x the (unique) label to x that, together with the labeling to the Y variables given by L', satisfies this constraint. So F_x contains α elements. Moreover, our assumption above implies that for at least $1-\beta$ of the variables $x \in X$, no element $i \in R$ appears more than $\gamma|F_x|$ times in F_x. By Claim 3.5, for such x, at least $1-\ell^2\gamma$ fraction of the variables in X' created from x have the property that it is impossible to satisfy more than one of their incident constraints simultaneously. Hence, the number of constraints in Φ' satisfied by L' is at most

$$\begin{aligned}
&\alpha^\ell \cdot \beta \cdot |X| \cdot \ell + \alpha^\ell (1-\beta)|X|\Big((1-\ell^2\gamma) + (\ell^2\gamma)\cdot \ell\Big) \\
&= |X'| \left(\beta\ell + (1-\beta)(1-\ell^2\gamma) + (1-\beta)(\ell^2\gamma)\ell\right) \\
&\le |E'| \left(\beta + \frac{1}{\ell} + (1-\beta)\ell^2\gamma\right).
\end{aligned}$$

■

The last lemma transforms a bipartite label cover into a non-bipartite label cover. We first prove a simple combinatorial claim.

Claim 3.7 *Let $A_1, \ldots, A_N$ be pairwise intersecting sets of size at most T. Then there exists an element contained in at least N/T of the sets.*

Proof: All sets intersect A_1 in at least one element. Since $|A_1| \le T$, there exists an element of A_1 contained in at least N/T of the sets. ■

For the following lemma, recall that a t-labeling labels each variable with a set of at most t labels. Recall also that a constraint on (x_1, x_2) is satisfied by a t-labeling L if there are labels $a \in L(x_1)$ and $b \in L(x_2)$ such that (a, b) satisfies the constraint.

Lemma 3.8 *There exists an efficient procedure that given an unweighted bipartite unique label cover instance $\Phi = (X, Y, R, \Psi, E)$ with all left-degrees equal to some ℓ, outputs an unweighted unique label cover instance $\Phi' = (X, R, \Psi', E')$ with the following properties:*

- *For any $\beta \ge 0$, if there exists a labeling L to Φ in which $1-\beta$ of the X variables have all their ℓ incident constraints satisfied, then there exists a labeling to Φ' and a set of $1-\beta$ of the variables of X such that all the constraints between them are satisfied.*
- *For any $\beta > 0$ and integer t, if there exists a t-labeling L' to Φ' and a set of β variables of X such that all the constraints between them are satisfied, then there exists a labeling L to Φ that satisfies at least β/t^2 of the constraints.*

Proof: For each pair of constraints $(x_1, y), (x_2, y) \in E$ that share a Y variable we add one constraint $(x_1, x_2) \in E'$. This constraint is satisfied when there exists a labeling to y that agrees with the labeling to x_1 and x_2. More precisely,

$$\psi'_{x_1x_2} = \Big\{(a_1, a_2) \in R \times R \;\Big|\; \exists b \in R \;(a_1, b) \in \psi_{x_1y} \wedge (a_2, b) \in \psi_{x_2y}\Big\}.$$

Notice that since the constraints in Ψ are unique, the constraints in Ψ' are also unique.

We now prove the first property. Let L be a labeling to Φ and let $C \subseteq X$ be of size $|C| \ge (1-\beta)|X|$ such that all constraints incident to variables in C are satisfied by L. Consider the labeling L' to Φ' given by $L'(x) = L(x)$. Then, we claim that L' satisfies all the constraints in Φ' between variables of C. Indeed, take any constraint between two variables $x_1, x_2 \in C$. Assume the constraint is created as a result of some $y \in Y$. Then, since $(L(x_1), L(y)) \in \psi_{x_1 y}$ and $(L(x_2), L(y)) \in \psi_{x_2 y}$, we also have $(L(x_1), L(x_2)) \in \psi'_{x_1 x_2}$.

It remains to prove the second property. Let L' be a t-labeling to Φ' and let $C \subseteq X$ be a set of variables of size $|C| \geq \beta |X|$ with the property that any constraint between variables of C is satisfied by L'. We first define a t-labeling L'' to Φ as follows. For each $x \in X$, we define $L''(x) = L(x)$. For each $y \in Y$, we define $L''(y) \in R$ as the label that maximizes the number of satisfied constraints between C and y. We claim that for each $y \in Y$, L'' satisfies at least $1/t$ of the constraints between C and y. Indeed, for each constraint between C and y consider the set of labels to y that satisfy it. These sets are pairwise intersecting since all constraints in Φ' between variables of C are satisfied by L'. Moreover, since Φ is a unique label cover, these sets are of size at most t. Claim 3.7 asserts the existence of a labeling to y that satisfies at least $1/t$ of the constraints between C and y. Since at least β of the constraints in Φ are incident to C, we obtain that L'' satisfies at least β/t of the constraints in Φ.

To complete the proof, we define a labeling L to Φ by $L(y) = L''(y)$ and $L(x)$ chosen uniformly from $L''(x)$. Since $|L''(x)| \leq t$ for all x, the expected number of satisfied constraints is at least β/t^2, as required. ■

4 Reduction to Vertex Cover in k-Uniform Hypergraphs

Throughout this section, we fix some $\varepsilon, \delta > 0$ and $k \geq 2$. The reader might wish to think of the case $k = 2$ at first reading. Our aim is to show a reduction from the problem described in Conjecture 3.1 to the Ek-Vertex-Cover problem. The vertices of the hypergraph we construct are weighted. One can obtain an unweighted hypergraph by using standard techniques (see, e.g., [10]). In the YES case, the hypergraph produced by the reduction contains an independent set of weight $1 - \frac{1}{k} - 2\varepsilon$ and in the NO case, the hypergraph contains no independent set of weight δ. It is easy to see that this implies the hardness of approximating Ek-Vertex-Cover to within any constant below k, assuming Conjecture 3.1.

4.1 Construction of the hypergraph

We define $p = 1 - \frac{1}{k} - \varepsilon$ as a bias parameter. The input to the reduction is a non-bipartite unweighted unique label cover instance $\Phi = (X, R, \Psi, E)$ as given by Conjecture 3.1 with parameters $\zeta = \varepsilon, \gamma = \delta/2$ and $t = t(k, \varepsilon, \delta)$ which will be chosen later. Notice that $|R|$ depends on ζ, γ and t and hence it is crucial that t does not depend on $|R|$.

The set of vertices is defined to be $X \times P(R)$. Hence, a vertex is a pair (x, F) where $x \in X$ is a variable and F is a subset of R. We define the *block* of a variable $x \in X$ as the set of vertices that correspond to x, i.e.,

$$B[x] \stackrel{def}{=} \{(x, F) \mid F \subseteq R\}.$$

The weight of a vertex (x, F) is defined to be

$$\frac{1}{|X|} \cdot \mu_p^R(F).$$

Thus the sum of the weights of all the vertices in the hypergraph equals 1.

Now we define the edges of the hypergraph. For any constraint $\psi_{x_1 x_2}$ in Ψ we define the following edges between the block $B[x_1]$ and the block $B[x_2]$:

$$\left\{ \{(x_1, G), (x_2, F_1), (x_2, F_2), \ldots, (x_2, F_{k-1})\} \;\middle|\; (G \times \cap_{i=1}^{k-1} F_i) \cap \psi_{x_1 x_2} = \emptyset \right\}.$$

In words, we create an edge $\{(x_1, G), (x_2, F_1), (x_2, F_2), \ldots, (x_2, F_{k-1})\}$ whenever there are no $a \in G$, $b \in \cap_{i=1}^{k-1} F_i$ that satisfy $(a, b) \in \psi_{x_1 x_2}$. Notice that every edge contains exactly k vertices, one vertex from the block $B[x_1]$ and $k-1$ vertices from the block $B[x_2]$. Also note that, as a result of parallel edges in E, we can have edges between $B[x_1]$ and $B[x_2]$ that correspond to more than one constraint.

4.2 YES case

Assume that Φ has a labeling L and a set of variables $X_0 \subseteq X$, $|X_0| \geq (1-\zeta)|X|$, such that all the constraints between variables in X_0 are satisfied by L. We claim that

$$\mathcal{IS} = \{(x, F) \mid x \in X_0,\ L(x) \in F\}$$

is an independent set. Consider any edge $\{(x_1, G), (x_2, F_1), \ldots, (x_2, F_{k-1})\}$ and let $\psi_{x_1 x_2}$ be the constraint it corresponds to. Assume on the contrary that all its vertices are in $\mathcal{IS}$. Clearly, this implies that $x_1 \in X_0$ and $x_2 \in X_0$. Hence, $\psi_{x_1 x_2}$ is satisfied by L and we have $(L(x_1), L(x_2)) \in \psi_{x_1 x_2}$. But since $L(x_1) \in G$ and $L(x_2) \in F_i$ for all $i = 1, \ldots, k-1$, this edge cannot exist in our hypergraph and we reach a contradiction.

To bound the weight of $\mathcal{IS}$, note that for every $x \in X_0$, $\mu_p^R(\mathcal{IS} \cap B[x]) = p$ (where we think of $\mathcal{IS} \cap B[x]$ as a subset of $P(R)$). Hence the weight of $\mathcal{IS}$ is

$$\frac{|X_0|}{|X|} \cdot p \geq (1-\zeta) \cdot \left(1 - \frac{1}{k} - \varepsilon\right) \geq 1 - \frac{1}{k} - 2\varepsilon$$

where we used $\zeta = \varepsilon$.

4.3 NO case

In this subsection we complete the proof by showing that if the hypergraph contains an independent set of weight δ then there exists a t-labeling L to Φ and a set $X^* \subseteq X$, $|X^*| \geq \gamma|X|$, such that all constraints between variables in X^* are satisfied by L.

So in the following, assume that the hypergraph contains an independent set $\mathcal{I}$ of weight δ. For every variable $x \in X$, let

$$\mathcal{F}[x] = \{F \subseteq R \mid (x, F) \in \mathcal{I}\}.$$

Our first observation is that we can assume without loss of generality that for all $x \in X$, $\mathcal{F}[x]$ is a monotone family. To see this, take any independent set $\mathcal{I}$ and for each $x \in X$ replace $\mathcal{I} \cap B[x]$ with its monotone extension (when considered as a subset of $P(R)$). Let $\mathcal{I}' \supseteq \mathcal{I}$ denote the resulting set. Clearly, its weight is at least that of $\mathcal{I}$. Moreover, we claim that $\mathcal{I}'$ is still an independent set. To prove this, assume on the contrary that there exists an edge $\{(x_1, G'), (x_2, F_1'), \ldots, (x_2, F_{k-1}')\}$ all of whose vertices are in $\mathcal{I}'$. By definition of $\mathcal{I}'$, there exist $G \subseteq G', F_1 \subseteq F_1', \ldots, F_{k-1} \subseteq F_{k-1}'$ such that all vertices of $\{(x_1, G), (x_2, F_1), \ldots, (x_2, F_{k-1})\}$ are in $\mathcal{I}$. Moreover, this tuple forms an edge since $(G \times \cap_{i=1}^{k-1} F_i) \cap \psi_{x_1 x_2} \subseteq (G' \times \cap_{i=1}^{k-1} F_i') \cap \psi_{x_1 x_2} = \emptyset$ and we reach a contradiction.

Let X^* be the set of variables $x \in X$ for which $\mu_p^R(\mathcal{F}[x]) \geq \delta/2$, i.e., a weight of at least $\delta/2$ of the total weight in the block $B[x]$ belongs to the independent set $\mathcal{I}$. By an averaging argument, we have $|X^*| \geq \delta|X|/2$. The next lemma completes the proof. It shows that there exists a t-labeling L that satisfies all the constraints between variables in X^*. Its proof is given in the next subsection.

Lemma 4.1 *Given $\mathcal{I}$ and X^* as above, there exists a $t = t(k, \varepsilon, \delta)$ and non-empty sets of labels $L[x] \subseteq R$ for every $x \in X^*$ such that*

- $\forall\ x \in X^*,\ |L[x]| \leq t$

- *For any constraint $\psi_{x_1x_2} \in \Psi$ with both $x_1, x_2 \in X^*$, there exist $a \in L[x_1]$, $b \in L[x_2]$, such that $(a, b) \in \psi_{x_1x_2}$.*

4.4 Proof of Lemma 4.1

We start with an overview of the proof. For each $x \in X^*$ we define the set $L[x]$ based on the family $\mathcal{F}[x]$. Since the family is monotone, it has low average sensitivity (after a slight shifting of the bias parameter) and hence it is approximated by a core-family $\widehat{\mathcal{F}}[x]$ with a core of "small" size. Since the core essentially captures the family $\mathcal{F}[x]$, it would be natural to define $L[x]$ to be the core. However, for the sake of analysis, we need to also include in $L[x]$ all elements that have non-negligible influence on the family $\mathcal{F}[x]$.[2] Recall that the average sensitivity is defined as the sum of the influences, and hence the number of elements that have non-negligible influence on $\mathcal{F}[x]$ is not too large. Thus the set $L[x]$ is not too large.

Next, we show that if $\psi_{x_1x_2} \in \Psi$ is any constraint with both $x_1, x_2 \in X^*$, then there exist $a \in L[x_1]$, $b \in L[x_2]$, such that $(a, b) \in \psi_{x_1x_2}$. To simplify the notation, we assume that the unique constraint $\psi_{x_1x_2}$ is of the form $\{(a, a) \mid a \in R\}$, i.e., the identity constraint. The proof for the general case is essentially identical and follows by applying to $\mathcal{F}[x_2]$ a permutation of the elements of R. With this assumption in place, our aim is to show that $L[x_1] \cap L[x_2] \neq \emptyset$.

Since $\mathcal{I}$ is an independent set, every set in $\mathcal{F}[x_1]$ and every $k-1$ sets in $\mathcal{F}[x_2]$ have a nonempty intersection. This seems to suggest that the cores of $\widehat{\mathcal{F}}[x_1]$ and of $\widehat{\mathcal{F}}[x_2]$ should have a nonempty intersection. To see why, notice that for any two nonempty core families $\mathcal{F}_1, \mathcal{F}_2$ with disjoint cores, one can find a set in $\mathcal{F}_1$ and a set in $\mathcal{F}_2$ that are disjoint. However, we are able to prove this only after including all influential elements as well, i.e., we show that $L[x_1] \cap L[x_2] \neq \emptyset$. We now proceed with the rather technical formal proof.

Let $\eta = \delta/(16k)$. Applying Theorem 2.6, we obtain

Lemma 4.2 *For every variable $x \in X^*$, there exists a real number $p[x] \in (1 - \frac{1}{k} - \varepsilon, 1 - \frac{1}{k} - \frac{\varepsilon}{2})$ and a core-family $\widehat{\mathcal{F}}[x] \subseteq P(R)$ with core $C[x]$ such that*

- *The average sensitivity* $\mathrm{as}_{p[x]}(\mathcal{F}[x]) \leq \frac{2}{\varepsilon}$.
- *The size of $C[x]$ is at most $t_0 = t_0(k, \varepsilon, \delta)$ (and crucially, t_0 is independent of $|R|$).*
- $\mu^R_{p[x]}(\mathcal{F}[x] \,\Delta\, \widehat{\mathcal{F}}[x]) < \eta$, *and in particular* $\mu^R_{p[x]}(\widehat{\mathcal{F}}[x]) \geq \delta/2 - \eta \geq \delta/4$.

Let $\eta' = \eta/(t_0 \cdot (2k)^{t_0})$ be a threshold parameter. For every $x \in X^*$, we define $\mathrm{Infl}[x] \subseteq R \setminus C[x]$ as the set of elements whose influence on the family $\mathcal{F}[x]$ is at least η', i.e.,

$$\mathrm{Infl}[x] = \{\sigma \in R \setminus C[x] \mid \mathrm{Inf}_{p[x]}(\mathcal{F}[x], \sigma) \geq \eta'\}.$$

Since $\mathcal{F}[x]$ has average sensitivity at most $\frac{2}{\varepsilon}$ and the average sensitivity is simply the sum of influences of all the elements, it follows that the size of $\mathrm{Infl}[x]$ is at most $\frac{2}{\eta'\varepsilon}$. Finally, we define the set $L[x]$ as

$$L[x] \stackrel{def}{=} C[x] \,\cup\, \mathrm{Infl}[x].$$

Clearly, $L[x]$ has size at most $t \stackrel{def}{=} t_0 + \frac{2}{\eta'\varepsilon}$. Notice that, as promised, t depends only on k, ε, δ and is independent of $|R|$.

[2]In [10], the union of the core and the set of all elements with non-negligible influence is referred to as the "extended core".

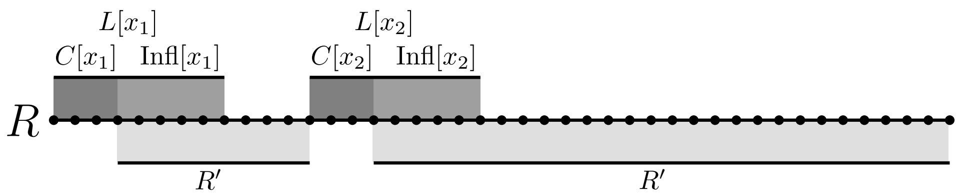

Figure 1: Subsets of R

Fix some constraint $\psi_{x_1x_2} \in \Psi$ with both $x_1, x_2 \in X^*$. To finish the proof of Lemma 4.1, it remains to show that there exist $a \in L[x_1]$, $b \in L[x_2]$, such that $(a,b) \in \psi_{x_1x_2}$. For simplicity, we assume in the following that the unique constraint $\psi_{x_1x_2}$ is of the form $\{(a,a) \mid a \in R\}$, i.e., the identity constraint. The proof for the general case is essentially identical and is left to the reader. With this assumption in place, our aim is to show that $L[x_1] \cap L[x_2] \neq \emptyset$.

Assume on the contrary that $L[x_1] \cap L[x_2] = \emptyset$. Our goal in the rest of the proof is to exhibit an edge $\{(x_1, G), (x_2, F_i)_{i=1}^{k-1}\}$ all of whose vertices are in the supposed independent set $\mathcal{I}$, thus giving a contradiction. We define R' as $R \setminus (C[x_1] \cup C[x_2])$. See Figure 1 for an illustration of the subsets of R that appear in this proof. We begin with a lemma.

Lemma 4.3 *There exists $U_0 \subseteq C[x_1]$ such that defining $\mathcal{H}[x_1] \subseteq P(R')$ as*

$$\mathcal{H}[x_1] \stackrel{def}{=} \{H \subseteq R' \mid U_0 \cup H \in \mathcal{F}[x_1]\ \}$$

we have $\mu_{p[x_1]}^{R'}(\mathcal{H}[x_1]) \geq 1 - 8\eta/\delta$.

Proof: The assumption $L[x_1] \cap L[x_2] = \emptyset$ implies in particular that $C[x_2] \cap L[x_1] = \emptyset$. Hence, every element of $C[x_2]$ has influence at most η' on the family $\mathcal{F}[x_1]$. Define $\mathcal{F}'[x_1] \subseteq \mathcal{F}[x_1]$ as

$$\mathcal{F}'[x_1] \stackrel{def}{=} \{F \in \mathcal{F}[x_1] \mid F \setminus C[x_2] \in \mathcal{F}[x_1]\}.$$

Then $\mathcal{F}'[x_1]$ is 'independent' of $C[x_2]$ in the sense that the containment of any F in $\mathcal{F}'[x_1]$ depends only on $F \setminus C[x_2]$. Applying Lemma 2.7, we get

$$\begin{aligned}
\mu_{p[x_1]}^{R}(\mathcal{F}[x_1] \setminus \mathcal{F}'[x_1]) &\leq \eta' \cdot |C[x_2]| \cdot (\min(p[x_1], 1 - p[x_1]))^{-|C[x_2]|} \\
&\leq \eta' \cdot t_0 \cdot (\min(p[x_1], 1 - p[x_1]))^{-t_0} \\
&\leq \eta' \cdot t_0 \cdot (2k)^{t_0} = \eta
\end{aligned}$$

by our choice of η'. It follows that

$$\mu_{p[x_1]}^{R}(\widehat{\mathcal{F}}[x_1] \ \setminus\ \mathcal{F}'[x_1]) \leq \mu_{p[x_1]}^{R}(\widehat{\mathcal{F}}[x_1] \ \setminus\ \mathcal{F}[x_1]) + \mu_{p[x_1]}^{R}(\mathcal{F}[x_1] \ \setminus\ \mathcal{F}'[x_1]) < 2\eta.$$

Intuitively, this says that except for some small measure 2η, $\widehat{\mathcal{F}}[x_1]$ is contained in $\mathcal{F}'[x_1]$. We use this, and the fact that the measure of the core family $\widehat{\mathcal{F}}[x_1]$ is at least $\delta/4 \gg 2\eta$ to conclude that there exists a set $U_0 \subseteq C[x_1]$ such that for almost all $D \subseteq R \setminus C[x_1]$ (under the measure $\mu_{p[x_1]}$) we have that $U_0 \cup D$ is in $\mathcal{F}'[x_1]$. More precisely,

$$\begin{aligned}
2\eta &> \mu_{p[x_1]}^{R}(\widehat{\mathcal{F}}[x_1] \ \setminus\ \mathcal{F}'[x_1]) \\
&= \sum_{U \subseteq C[x_1]} \mu_{p[x_1]}^{R}(\{D \subseteq R \mid D \cap C[x_1] = U \text{ and } D \in \widehat{\mathcal{F}}[x_1] \ \setminus\ \mathcal{F}'[x_1]\}) \\
&= \sum_{U \subseteq C[x_1],\ U \in \widehat{\mathcal{F}}[x_1]} \mu_{p[x_1]}^{C[x_1]}(U) \cdot \mu_{p[x_1]}^{R \setminus C[x_1]}(\{D \subseteq R \setminus C[x_1] \mid (U \cup D) \notin \mathcal{F}'[x_1]\})
\end{aligned}$$

where the last equality holds since the condition $D \in \widehat{\mathcal{F}}[x_1]$ is equivalent to $D \cap C[x_1] \in \widehat{\mathcal{F}}[x_1]$ and since μ_p is a product measure. Together with

$$\sum_{U \subseteq C[x_1],\ U \in \widehat{\mathcal{F}}[x_1]} \mu_{p[x_1]}^{C[x_1]}(U) = \mu_{p[x_1]}^{R}(\widehat{\mathcal{F}}[x_1]) \geq \delta/4$$

we obtain that there exists a $U_0 \subseteq C[x_1]$, $U_0 \in \widehat{\mathcal{F}}[x_1]$ for which

$$\mu_{p[x_1]}^{R \setminus C[x_1]}(\{D \subseteq R \setminus C[x_1] \mid U_0 \cup D \notin \mathcal{F}'[x_1]\}) < 2\eta/(\delta/4).$$

Equivalently, we have

$$\mu_{p[x_1]}^{R \setminus C[x_1]}(\{D \subseteq R \setminus C[x_1] \mid U_0 \cup D \in \mathcal{F}'[x_1]\}) \geq 1 - 8\eta/\delta.$$

Since $\mathcal{F}'[x_1]$ is independent of $C[x_2]$ (in the sense described above), this inequality is equivalent to

$$\mu_{p[x_1]}^{R'}(\{H \subseteq R' \mid U_0 \cup H \in \mathcal{F}'[x_1]\}) \geq 1 - 8\eta/\delta.$$

It remains to recall that $\mathcal{F}'[x_1] \subseteq \mathcal{F}[x_1]$. ∎

Analogous to Lemma 4.3 we have by symmetry,

Lemma 4.4 *There exist $V_0 \subseteq C[x_2]$ such that defining $\mathcal{H}[x_2] \subseteq P(R')$ as*

$$\mathcal{H}[x_2] \stackrel{def}{=} \{H \subseteq R' \mid V_0 \cup H \in \mathcal{F}[x_2]\ \}$$

we have $\mu_{p[x_2]}^{R'}(\mathcal{H}[x_2]) \geq 1 - 8\eta/\delta$.

Let $p^* \stackrel{def}{=} 1 - \frac{1}{k} - \frac{\varepsilon}{2}$. Since $\mathcal{F}[x_1]$ and $\mathcal{F}[x_2]$ were assumed to be monotone, we have that $\mathcal{H}[x_1]$ and $\mathcal{H}[x_2]$ are monotone subfamilies of $P(R')$. Therefore, according to Theorem 2.3, $\mu_{p^*}^{R'}(\mathcal{H}[x_1]) \geq \mu_{p[x_1]}^{R'}(\mathcal{H}[x_1]) \geq 1 - 8\eta/\delta$ and similarly for x_2. Hence, the intersection of the families $\mathcal{H}[x_1]$ and $\mathcal{H}[x_2]$ satisfies

$$\mu_{p^*}^{R'}(\mathcal{H}[x_1] \cap \mathcal{H}[x_2]) \geq 1 - 16\eta/\delta = 1 - \frac{1}{k}$$

by our choice of η. Hence, Lemma 2.9 implies that there exist sets $H_1, H_2, \ldots, H_k \in \mathcal{H}[x_1] \cap \mathcal{H}[x_2]$ such that

$$\cap_{i=1}^{k} H_i = \emptyset.$$

In particular, $H_1, H_2, \ldots, H_{k-1} \in \mathcal{H}[x_2]$ and $H_k \in \mathcal{H}[x_1]$.

Now define $G = U_0 \cup H_k$ and $F_i = V_0 \cup H_i$ for $1 \leq i \leq k-1$. By definition of the families $\mathcal{H}[x_1], \mathcal{H}[x_2]$, we have $G \in \mathcal{F}[x_1], F_i \in \mathcal{F}[x_2]$ for $1 \leq i \leq k-1$. Thus $\{(x_1, G), (x_2, F_i)_{i=1}^{k-1}\}$ are vertices in the supposed independent set and they form an edge since

$$G \cap \left(\cap_{i=1}^{k-1} F_i \right) = \cap_{i=1}^{k} H_i = \emptyset.$$

This completes the proof.

References

[1] S. Arora, B. Bollobás, and L. Lovász. Proving integrality gaps without knowing the linear program. In *Proceedings of the 43rd Annual Symposium on Foundations of Computer Science (FOCS) 2002*, pages 313–322, Vancouver, Canada, Nov. 2002.

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

[3] S. Arora and S. Safra. Probabilistic checking of proofs: a new characterization of NP. *J. ACM*, 45(1):70–122, 1998.

[4] M. Bellare, O. Goldreich, and M. Sudan. Free bits, PCPs, and nonapproximability—towards tight results. *SIAM Journal on Computing*, 27(3):804–915, June 1998.

[5] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for unique games. In *Proceedings of the 38th ACM Symposium on Theory of Computing (STOC)*, 2006.

[6] S. Chawla, R. Krauthgamer, R. Kumar, Y. Rabani, and D. Sivakumar. On the hardness of approximating sparsest cut and multicut. In *Proc. of 20th IEEE Annual Conference on Computational Complexity (CCC)*, pages 144–153, 2005.

[7] I. Dinur, V. Guruswami, and S. Khot. Vertex cover on k-uniform hypergraphs is hard to approximate within factor $(k-3-\varepsilon)$. In *Electronic Colloquium on Computational Complexity, Technical Report TR02-027*, 2002.

[8] I. Dinur, V. Guruswami, S. Khot, and O. Regev. A new multilayered PCP and the hardness of hypergraph vertex cover. *SIAM Journal on Computing*, 34(5):1129–1146, 2005.

[9] I. Dinur, E. Mossel, and O. Regev. Conditional hardness for approximate coloring. In *Proc. 38th ACM Symp. on Theory of Computing (STOC)*, 2006.

[10] I. Dinur and S. Safra. On the hardness of approximating minimum vertex cover. *Annals of Mathematics*, 162(1), 2005. Preliminary version in STOC 2002.

[11] U. Feige, S. Goldwasser, L. Lovász, S. Safra, and M. Szegedy. Approximating clique is almost NP-complete. In *Proc. 32nd IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 2–12, 1991.

[12] E. Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):27–35, 1998.

[13] O. Goldreich. Using the FGLSS-reduction to prove inapproximability results for minimum vertex cover in hypergraphs. ECCC Technical Report TR01-102, 2001.

[14] A. Gupta and K. Talwar. Approximating unique games. In *Proceedings of Annual Symposium on Discrete Algorithms (SODA)*, pages 99–106, 2006.

[15] V. Guruswami, J. Håstad, and M. Sudan. Hardness of approximate hypergraph coloring. *SIAM J. Comput.*, 31(6):1663–1686, 2002.

[16] E. Halperin. Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs. *SIAM Journal on Computing*, 31(5):1608–1623, 2002.

[17] J. Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. *Acta Math.*, 182(1):105–142, 1999.

[18] J. Håstad. Some optimal inapproximability results. *Journal of ACM*, 48(4):798–859, 2001.

[19] J. Holmerin. Improved inapproximability results for vertex cover on k-regular hyper-graphs. In *International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 1005–1016, 2002.

[20] J. Holmerin. Vertex cover on 4-regular hyper-graphs is hard to approximate within $2 - \varepsilon$. In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 544–552, 2002.

[21] G. Karakostas. A better approximation ratio for the vertex cover problem. In *Proceedings of the 32nd International Colloquium on Automata, Languages and Programming (ICALP)*, pages 1043–1050, Lisboa, Portugal, July 2005.

[22] S. Khot. On the power of unique 2-Prover 1-Round games. In *Proc. 34th ACM Symp. on Theory of Computing (STOC)*, pages 767–775, May 2002.

[23] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? In *Proc. 45th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 146–154, 2004.

[24] S. Khot and N. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *Proc. 46th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 53–62, 2005.

[25] G. Margulis. Probabilistic characteristics of graphs with large connectivity. *Prob. Peredachi Inform.*, 10:101–108, 1974.

[26] E. Mossel, R. O'Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences, invariance and optimality. In *Proc. 46th IEEE Symp. on Foundations of Computer Science (FOCS)*, pages 21–30, 2005.

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

[28] L. Russo. An approximate zero-one law. *Z. Wahrsch. Verw. Gebiete*, 61(1):129–139, 1982.

[29] L. Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *Proc. 33rd ACM Symp. on Theory of Computing (STOC)*, pages 453–461, 2001.

[30] L. Trevisan. Approximation algorithms for unique games. In *Proceedings of the 46th Annual Symposium on Foundations of Computer Science (FOCS) 2005*, pages 197–205, 2005.