Covering CSPs

Irit Dinur [†]
Weizmann Institute

Gillat Kol [§]
Weizmann Institute

***Abstract*—We study the covering complexity of constraint satisfaction problems (CSPs). The covering number of a CSP instance C, denoted v(C), is the smallest number of assignments to the variables, such that each constraint is satisfied by at least one of the assignments. This covering notion describes situations in which we must satisfy all the constraints, and are willing to use more than one assignment to do so. At the same time, we want to minimize the number of assignments.**

We study the covering problem for different constraint predicates. We first observe that if the predicate contains an odd predicate, then it is covered by any assignment and its negation. In particular, 3CNF and 3LIN, that are hard in the max-CSP sense, are easy to cover. However, the covering problem is hard for predicates that do not contain an odd predicate:

1. For the 4LIN predicate, it is NP-hard to decide if a given instance C has v(C) at most 2, or v(C) is super-constant.

2. (a) We propose a framework of covering dictatorship tests. We design and analyze such a dictatorship test for every predicate that supports a pairwise independent distribution. (b) We introduce a covering unique games conjecture, and use it to convert the covering dictatorship tests into conditional hardness results.

3. Finally, we study a hypothesis about the hardness of covering random instances that is similar to Feige's R3SAT hypothesis. We show the following somewhat surprising implication: If our hypothesis holds for dense enough instances, then it is hard to color an O(1)-colorable hypergraph with a polynomial number of colors.

I. Introduction

We study the *covering complexity* of constraint satisfaction problems (CSPs). Let φ be a predicate, and let $\mathcal{C}$ be a φ-CSP instance, which is a set of φ-constraints over n boolean variables and their negations. The *covering number* of $\mathcal{C}$, denoted $\nu(\mathcal{C})$, is the smallest number of assignments to the variables that "*covers*" all of the constraints, i.e., such that each constraint is satisfied by at least one of the assignments. We denote by cover-φ the problem of finding the covering number of a given φ-CSP instance.

The notion of cover-CSPs differs from the standard notion of max-CSPs, as they each operate under a different restriction and try optimize a different aspect of the problem given the restriction: The notion of max-CSPs is relevant when we restrict ourselves to a single assignment and want to maximize the fraction of satisfied constraints. In contrast, the notion of a covering number is of interest when we must satisfy all or nearly all of the constraints, and are willing to use more than one assignment to do so. Our goal is then to minimize the number of needed solutions.

One example of a situation described by the covering number is the dinner party problem: You are having some friends over for dinner, and each one has different dietary constraints. You want everyone to have at least something to eat, and at the same time would like to cook as few dishes as possible. Another example is when designing a system of health care centers, each offering different services, that will be accessible and will meet the needs of all patients.

Finding the exact covering number is $\mathcal{NP}$-hard for many interesting predicates φ. Therefore, we study the hardness of approximating this value, namely minimizing the number of solutions that together cover all of the constraints. Formally, we define the following gap problem:

gap-cover-$\varphi_{c,s}$ problem: Let $c < s \in \mathbb{N}$. Given a φ-CSP instance $\mathcal{C}$, decide between

- **Yes case:** $\nu(\mathcal{C}) \leq c$.
- **No case:** $\nu(\mathcal{C}) \geq s$.

As is done for the max-CSP case, we study the covering problem for different predicates φ, and seek a characterization of predicates that are covering-hard to approximate. It turns out that the set of predicates which are covering hard to approximate is very different from the set of predicates that are hard to approximate in the max-CSP sense. In fact we show that the sets are (in a sense) incomparable.

Covering and Coloring: Covering CSPs can be viewed as a generalization of graph (or hypergraph) coloring problems. A coloring problem is given by a system of not-equal (or not-all-equal) constraints on a set of vertices. It has already been observed by [7] that a graph (hypergraph) is 2^c colorable iff there are c assignments to the variables that cover all constraints. Our new notion of covering CSPs extends that of coloring as follows. It is natural to allow an algorithm "more" colors when attempting to legally color a graph, yet, in contrast, it is usually meaningless to allow "more" alphabet symbols for satisfying a φ-CSP for a general predicate φ. The covering formulation gives a

[†]Research supported in part by the Israel Science Foundation grant no. 1179/09 and by the Binational Science Foundation grant no. 2008293 and by an ERC grant no. 239985. Part of the research done while visiting Microsoft Research in New England.

[§]Research supported by the Israel Science Foundation and by the NSF. Part of the research done while visiting Microsoft Research in New England.

natural way in which "more colors" can be used in satisfying a φ-CSP for any φ.

We mention that the paper [7] introduces the related notion of "covering PCPs" and proves hardness of approximate hypergraph coloring by analyzing the hardness of covering the not-all-equal predicate. Interestingly, our work reveals that understanding the hardness of covering the not-all-equal predicate is central for any covering-CSP problem.

A. Our Results

We first observe that odd predicates φ (i.e., predicates $\varphi : \{\pm 1\}^t \to \{\pm 1\}$ for which $\forall x : \varphi(x) = -\varphi(-x)$) are easy to cover: Any pair of an assignment a and its negation $-a$ will cover the entire instance, since always either a or $-a$ causes φ to be true. Moreover, let $\mathcal{O}$ be the set of predicates φ containing an odd predicate, all the predicates $\varphi \in \mathcal{O}$ are easy. Formally, we define $\mathcal{O}$ to be the set of predicates $\varphi : \{\pm 1\}^t \to \{\pm 1\}$ (as is customary, we view a $(-1) = (-1)^1$ value as "true") satisfying:

$$\mathcal{O} = \left\{ \varphi \,\middle|\, \forall x \in \{\pm 1\}^t : \varphi(x) = -1 \text{ or } \varphi(-x) = -1 \right\}.$$

Observation. *Let $\varphi \in \mathcal{O}$, and let $\mathcal{C}$ be a φ-CSP instance. Then $\nu(\mathcal{C}) \leq 2$.*

In particular, 3CNF and 3LIN which are both very hard to approximate in the max-CSP sense, are easy in the covering sense.

1) Covering Hardness of 4LIN*:* In contrast to 3LIN, we show that the predicate $\varphi = \mathsf{4LIN}$, that only accepts inputs with an odd number of 1s, is $\mathcal{NP}$-hard. Formally, for $\mathsf{4LIN} : \{\pm 1\}^4 \to \{\pm 1\}$, $\mathsf{4LIN}(x_1, x_2, x_3, x_4) = x_1 x_2 x_3 x_4$, we show:

Theorem 1. *gap-cover-$\mathsf{4LIN}_{2,k}$ is $\mathcal{NP}$-hard for every $k \in \mathbb{N}$. Furthermore, for sufficiently small $\epsilon > 0$, the following holds: In the yes case the instance is coverable by two assignments, each of which (seperatly) satisfies $1-\epsilon$ fraction of the constraints. In the no case, no k assignments cover more than $1 - \frac{1}{2^k} + 20\sqrt{\epsilon}$ fraction of the constraints.*

Observe that the problem gap-cover-$\mathsf{4LIN}_{1,k}$ is easy for every k, as we can run a Gaussian elimination process to check whether there exists a single assignment that satisfies all the constraints.

We mention that our result can be viewed as a strengthening of Håstad's hardness result [8] for linear predicates with even arity≥ 4, since in the yes case there is a solution that satisfies $1 - \epsilon$ fraction of the constraints (actually, there are at least two such solutions). Furthermore, observe that k random assignments are expected to satisfy $1 - \frac{1}{2^k}$ fraction of the constraints, thus the no case shows that no k assignments can cover significantly more constraints than random k assignments.

Our proof of Theorem 1 relies on a dictatorship test whose analysis extends the analysis of [7] of the hardness of covering the 4-not-all-equal predicate, using the language of the invariance principle developed by [15], [14], [17], [16].

2) Characterization of Covering-Hard Predicates: We conjecture that for every $\varphi \notin \mathcal{O}$ the cover-φ problem is hard to approximate, and are able to partially prove this conjecture. To do so, we offer a φ-based covering dictatorship test. We then suggest a covering conjecture that corresponds to the unique games conjecture, and show how to use the dictatorship test to obtain conditional hardness of covering results.

Covering dictatorship test for a general predicate φ: We develop a general framework for covering dictatorship tests using a given predicate φ. The completeness and soundness criteria are different in the covering world:

- In the yes case, two dictators perfectly cover the test's constraints.
- In the no case, any *regular* set of functions $F = \{f_1, \ldots, f_k\}$ fails to cover all of the test's constraints. F is called regular if for any $K \subseteq [k]$ the product function $f_K = \prod_{\ell \in K} f_\ell$ is far from a dictatorial function (i.e., all its influences are low). We mention that this involved soundness condition is inherent, see Section I-B2.

Following Austrin and Mossel [2] we prove

Theorem 2. *Let $\varphi \notin \mathcal{O}$, and assume that there exists a balanced, pairwise independent distribution on the support of φ. Then there exists a φ-based covering-dictatorship test with completeness 2 and soundness k, for every $k \in \mathbb{N}$.*

We remark that every predicate $\varphi \notin \mathcal{O}$ that does not have degree-1 and degree-2 terms in its Fourier expansion, satisfies the condition of the theorem.

Covering unique games hardness for a general predicate φ: We suggest the following covering conjecture that corresponds to the unique games conjecture:

Conjecture 3 (**Covering Unique Games**)**.** *There exists $c \in \mathbb{N}$ such that for every sufficiently small $\delta > 0$ there exists $R \in \mathbb{N}$ such that the following holds. Given a bipartite label cover instance $\mathcal{LC}$ with permutation constraints over label set $[R]$ and vertex set $U \times V$, it is $\mathcal{NP}$-hard to decide between:*

- ***Yes case:** There exist c assignments such that for every vertex $u \in U$, at least one of the assignments satisfies all the edges touching u.*
- ***No case:** $\mathsf{OPT}(\mathcal{LC}) \leq \delta$. I.e., every assignment satisfies at most δ fraction of the edge constraints.*

We mention that Khot and Regev [12] consider a similar conjecture in the max-CSP setting: In the yes case they require a single assignment that for $1 - \delta$ fraction of the vertices $u \in U$, satisfies all the edges touching u. They show that their conjecture is equivalent to the unique games conjecture. See further discussion regarding the formulation of our covering conjecture in Section V-A.

Our conjecture is clearly false with $c = 1$, but as far as we know may be true with even $c = 2$. The conjecture is incomparable to the unique games conjecture (our completeness does not require any single assignment to satisfy a large fraction of edges). However it clearly implies the unique games conjecture with completeness $\frac{1}{c}$ (instead of $1-\epsilon$).

As usual, we say that a problem P is covering unique games hard, if it is hard assuming Conjecture 3.

Theorem 4. *Let $\varphi \notin \mathcal{O}$, and assume that there exists a balanced, pairwise independent distribution on the support of φ. Let c be the completeness constant from the covering unique games conjecture. Then gap-cover-$\varphi_{2c,k}$ is covering unique games-hard for every $k \in \mathbb{N}$.*

3) Hardness of Approximate Coloring and Covering Random CSP Instances: We now return to the problem of approximate coloring: Given an $O(1)$-colorable graph (or hypergraph), what is the smallest number of colors needed to color it in polynomial time? This is a notorious open question with an exponential gap between known upper [9], [3], [1] and lower [7], [10], [4] bounds. One might hope that viewing this classical problem in the broader context of covering-CSPs may shed new light on it.

We show some progress in this direction, if one is willing to assume hardness of covering *random* CSP instances.

In a seminal paper, Feige [5] hypothesizes that no polynomial time algorithm is able to distinguish between a random 3SAT and a satisfiable one, and shows that this implies various hardness of approximation results. However, since a 3SAT instance is always coverable by 2 assignments, it seems impossible to derive a hardness of coloring results from Feige's hypothesis.

We formulate an analogous hypothesis about the hardness of distinguishing between random and 2-coverable 4LIN-CSP instances. We prove that if our hypothesis holds with sufficient density, it implies hardness of approximate hypergraph coloring to within *polynomial* factors.

Hypothesis 5 (Covering 4LIN Hypothesis, with density parameter Δ). *There is no polynomial time algorithm that outputs* `typical` *for most 4LIN-CSP instances with n variables and $m = \Delta \cdot n$ clauses, and never outputs* `typical` *for a 2-coverable 4LIN-CSP instance.*

We point out that our $\mathcal{NP}$-hardness result (Theorem 1) implies that none of the currently known algorithmic techniques can refute this hypothesis. Furthermore, the best known algorithms can only refute instances with density at least $\Delta \geq n^{0.5}$ [6].

Theorem 6. *If Hypothesis 5 holds with density parameter $\Delta = n^{\delta}$ for some positive $\delta > 0$, then it is hard to decide if a 4-uniform hypergraph is 4-colorable or requires at least a polynomial number of colors.*

In Section VI we formulate a (weaker) hypothesis for covering a general predicate φ, and show that it has the same implication.

B. Technique

We face two main challenges: The first is achieving *perfect* covering completeness (being able to cover *all* the constraints vs. covering $1-\epsilon$ of them). We introduce a technique of "duplicating" the label cover instance and design an appropriate correlated-noise dictatorship test. The basic technique is explained below, variations of it are used in the first two parts of the work. The second challenge is in handling several assignments at once when proving the soundness property. This involves solving several different problems, some of which are very roughly described below. We next present a very informal discussion of our efforts.

1) Achieving perfect completeness: Hardness of approximation results for CSPs are usually obtained through a dictatorship test for a given function $f : \{\pm 1\}^R \to \{\pm 1\}$. A typical dictatorship test involves selecting a few points in $\{\pm 1\}^R$ and then querying the function f on slight perturbations of these points. The perturbation usually involves flipping the value of each coordinate independently with small probability ϵ. While the perturbation is very effective in killing the large Fourier coefficients of f, it also "ruins" the perfect completeness, causing even a perfect dictator to be accepted with probability $1-\epsilon$.

To overcome this problem we offer the new notion of a *duplicated* label cover instance: Given a label cover instance, each constraint $\pi_{v,u} : [R] \to [R]$ will be extended to the "duplicated" constraint $\pi_{v,u} : [2R] \to [2R]$ by

$$\forall j \in [R], \quad \pi_{v,u}(j+R) = \pi_{v,u}(j) + R.$$

This notion of a duplicated label cover will be central in our work. Observe that if $L : V \to [R]$ satisfies the constraints in the original label cover, then both L and $L+R$ satisfy the constraints in the duplicated label cover. This allows us to design a dictatorship test with enough random noise to eliminate the large Fourier coefficients, without hurting the perfect completeness. The idea is that independently for each pair of coordinates $j, j+R$, noise will be applied to at most one of the two coordinates.

2) Dealing with several proofs: When proving covering soundness in a dictatorship test we have to analyze the test's behavior on several functions at once, which means an involved rejection probability expression. This expression is basically the product of the expressions for the individual functions.

One complication arises from the fact that the test might be completely covered even if none of the functions are "dictatorial". For example, suppose that f is a random function and $f' = f \cdot x_j$. Then always either $f(x)f(y)f(z)f(-xyz) = -1$ or $f'(x)f'(y)f'(z)f'(-xyz) = -1$. This means that the natural 4LIN test will always pass while both f and f' are completely random functions. The reason this happens is because $f \cdot f'$ is a dictatorship, forcing our analysis to consider all possible products of the given functions.

This brings about another complication, which is that even if all given functions are "folded", or balanced, their product does not have to be. This means that the empty Fourier coefficient may be large, which complicates the analysis.

Covering hardness of 4LIN*:* Both of the above problems were faced by [7] when analyzing the covering soundness of the NAE_4 predicate. The technique of [7] does carry over (with some adaptation originating from our correlated noise) to proving covering soundness for the 4LIN predicate. Indeed, we use a test very similar to theirs (this test was originally suggested by Håstad [8]). However, we analyze the test in the more recent framework of the invariance principle developed by [15], [14]. This technique follows our intuition of the problem better, and is less "tailor-made" for specific predicates (indeed, in the second part of the work we use the invariance principle to show a more general result).

We mention that we cannot use the invariance principle directly, and that the usage of the invariance principle to obtain $\mathcal{NP}$-hardness results (as opposed to conditional results) is challenging. Similar difficulties were recently faced by [17], [16], and we indeed use parts of their analysis.

Covering dictatorship test for a general predicate φ: Our starting point for analyzing a general predicate φ is the work of [2], who considered any predicate φ that contains a pairwise independent distribution in its support. Their test uses independent noise, has a simpler rejection expression, and also assumes folding.

To analyze our test we rely on the following observation: every predicate $\varphi \notin \mathcal{O}$ is contained in a shifted NAE predicate (see Claim II.1). Equipped with this observation, we bound the rejection term by exploiting the symmetry of the Fourier expansion of NAE, extending a 'pairing' trick from [7]. We mention that we cannot simply reduce the covering soundness of NAE_4 to that of φ, as the distribution used by our dictatorship test must be supported on φ for maintaining completeness. Still, as it turns out, the key for analyzing the rejection term for a general predicate $\varphi \notin \mathcal{O}$ is analyzing the same term for NAE.

Covering unique games hardness for a general predicate φ: Having developed a dictatorship test, the "straightforward" path, following [11], is to analyze it for a function f that is the average of the long-code functions $f_v : \{\pm 1\}^R \to \{\pm 1\}$ for all neighbors v of a given u. An influential coordinate for f implies a consistent influential coordinate for many f_v's.

In our case, however, since we have k proofs, we also have k expected functions $f_1, \ldots, f_k$ for the same vertex u. If these k functions cover the dictatorship test we can only deduce that there is a product of $f_1, \ldots, f_k$ that has an influential coordinate. Constructing a good assignment for the label cover instance becomes non-trivial. To solve this we must analyze the dictatorship test in the more general multi-function setting.

C. Future Work

This work is a first step in studying the covering complexity of CSPs, and there are many interesting directions to pursue. Aside from the obvious directions of proving more $\mathcal{NP}$-hardness results, we mention a couple of directions:

Quantitative results: Our NP-hardness result for 4LIN implies hardness of covering a 2-coverable instance with about $\Omega(\log \log \log n)$ assignments. For the special case of coloring, better quantitative results are known, corresponding to a gap between $O(1)$ and $\Omega(\log \log n)$ [10]. On the other hand, the best algorithms require $O(\log n)$ assignments corresponding to a polynomial number of colors. Our question is: Is there any predicate φ for which one can prove $\mathcal{NP}$-hardness with a gap of $O(1)$ and $\Omega(\log n)$? In fact, any gap between $O(1)$ and $\omega(\log \log n)$ would be interesting.

More general characterization: What is the complexity of covering predicates $\varphi \notin \mathcal{O}$ that do not support a pairwise independent distribution?

Covering unique games conjecture: What can be said about the covering unique games conjecture? Can it be related to other conjectures such as Khot's d-to-1 conjecture?

Reductions between covering problems: Can one devise 'direct' reductions between covering problems? For example, does cover-NAE reduce to cover-φ for some other predicate φ? Gadget reductions simply fail in this context, and it would be interesting to find alternatives.

D. Organization

We begin in Section II with preliminaries and definitions. The covering hardness of 4LIN is proved in Section III. The characterization of covering-hard predicates can be found in Sections IV-V, where Section IV is devoted to the covering dictatorship test, and Section V is devoted to the covering unique games hardness result. Finally, the relations between random CSP instances and hardness of approximate coloring are discussed in Section VI.

II. Definitions and Preliminaries

A. Covering Problems

Let $X = \{x_1, ..., x_n\}$ be a set of n boolean variables, each taking a value in $\{\pm 1\}$. As is customary, we view a $(-1) = (-1)^1$ value as "*true*", and a $1 = (-1)^0$ value as "*false*" (e.g., $1 \wedge (-1) = 1$). Let $\varphi : \{\pm 1\}^t \to \{\pm 1\}$ be a predicate. A *φ-constraint* over X is an equation of the form $\varphi(\sigma_1 x_{i_1}, \ldots, \sigma_t x_{i_t}) = b$, where $i_1, \ldots, i_t \in [n]$ and $b, \sigma_1, \ldots, \sigma_t \in \{\pm 1\}$. A *$\varphi$-CSP instance* $\mathcal{C}$ is a set of φ-constraints over X.

Let $\mathcal{L} \subseteq \{\pm 1\}^n$ be a set of assignments for X. We say that $\mathcal{L}$ *covers* the instance $\mathcal{C}$ if for every constraint in $\mathcal{C}$, there exists an assignment in $\mathcal{L}$ that satisfies it. The *covering number* of $\mathcal{C}$, denoted $\nu(\mathcal{C})$, is the smallest number of assignments for X such that each constraint is satisfied by at least one of the assignments. We denote by cover-φ the problem of finding the covering number of a given CSP.

1) Containment in NAE*:* The following claim shows that the support of any predicated $\varphi \notin \mathcal{O}$ is contained in the support of NAE, upto a "sign". The claim will be very useful to us, as it allows us to move from a general predicate φ to the specific predicate NAE. Recall -1 denotes acceptance:

Claim II.1. *For every $\varphi \notin \mathcal{O}$, $\varphi : \{\pm 1\}^t \to \{\pm 1\}$, there is a "sign" $\sigma = (\sigma_1, \ldots, \sigma_t) \in \{\pm 1\}^t$ such that $\forall x \in \{\pm 1\}^t$: $\varphi(\sigma_1 x_1, \ldots, \sigma_t x_t) \geq \mathsf{NAE}_t(x_1, \ldots, x_t)$.*

The claim easily follows from the fact that for a predicate $\varphi \notin \mathcal{O}$ there exists an assignment a and its negation $-a$ that are both rejected by φ. Thus, by taking $\sigma = a$ we get that $\varphi(\sigma_1 x_1, \ldots, \sigma_t x_t)$ rejects both the assignment 1^t and $(-1)^t$, and thus its support is contained in the support of NAE_t.

B. Label Cover

A *bipartite label cover instance* is a tuple $\mathcal{LC} = (U, V, E, R_1, R_2, \Pi)$. Here U and V are the two vertex sets of a bipartite graph, and E is the set of edges between U and V. R_1 and R_2 satisfy $R_1 \leq R_2 \in \mathbb{N}$. $[R_1]$ is the set of labels for vertices in U, and $[R_2]$ is the set of labels for vertices in V. Π is a collection of "*projections*", one for each edge in E. That is, $\Pi = \{\pi_{v,u} : [R_2] \to [R_1]\}_{(u,v)\in E}$.

Let L be an assignment for the vertices in $U \times V$, that assigns to each vertex in U a label from $[R_1]$, and to each vertex in V a label from $[R_2]$. Let $(u, v) \in E$ be an edge. We say that L *satisfies* the edge (u, v) if $\pi_{v,u}(L(v)) = L(u)$. The value of a label cover instance $\mathcal{LC}$, denoted $\mathsf{OPT}(\mathcal{LC})$, is the maximal fraction of satisfied edges over all assignments L. It is well know that it is $\mathcal{NP}$-hard to approximate the value of a given label cover instance.

1) Smooth Label Cover: A *smooth* label cover instance is a label cover instance that satisfies the following: Let $v \in V$. In expectation over neighbors u of v, every large set of assignments for v induces a large set of assignments for u. In other words, for a sufficiently large $A \subseteq [R_2]$, it holds that $\mathbb{E}_{u\in\Gamma(v)} |\pi_{v,u}(A)|$ is large.

Lemma II.2 ([13])**.** *Let $\epsilon > 0$ be a sufficiently small constant and let $r \in \mathbb{N}$ be a sufficiently large constant. There exists an efficient transformation that maps an instance ψ of* $\mathsf{3SAT}$ *to an instance $\mathcal{LC}_{\epsilon,r} = (U, V, E, R_1, R_2, \Pi)$ of bipartite label cover such that*

- ***Completeness:*** *If ψ is satisfiable, $\mathsf{OPT}(\mathcal{LC}_{\epsilon,r}) = 1$.*
- ***Soundness:*** *If ψ is unsatisfiable, $\mathsf{OPT}(\mathcal{LC}_{\epsilon,r}) < c_0^r$, where $c_0 \in (0, 1)$ is an absolute constant.*
- ***Smoothness:*** *For every vertex $v \in V$ and any subset of labels $A \subseteq [R_2]$ satisfying $|A| \geq \frac{1}{\epsilon^3}$, it holds that*

$$\Pr_{u\in\Gamma(v)}\left[|\pi_{v,u}(A)| \geq \frac{1}{\epsilon^2}\right] \geq 1 - 2\epsilon.$$

2) Label Cover with Permutation Constraints: Of particular interest to us are bipartite label cover instances with *permutation constraints*. Namely, where $R_1 = R_2 = R \in \mathbb{N}$ (that is, the sets of labels for U and V are the same), and $\Pi = \{\pi_{v,u} : [R] \to [R]\}_{(u,v)\in E}$ is a collection of *permutations.*

3) Duplicated Label Cover: We define the new notion of a *duplicated* label cover instance, which will play a main role in our proofs. We assume to be given a bipartite label cover instance $\mathcal{LC}' = (U, V, E, R_1, R_2, \Pi')$ with $\Pi' = \left\{\pi'_{v,u} : [R_2] \to [R_1]\right\}_{(u,v)\in E}$. The *duplicated-$\mathcal{LC}'$* instance is a new bipartite label cover instance $\mathcal{LC} = (U, V, E, R_1, 2R_2, \Pi)$, where $\Pi = \{\pi_{v,u} : [2R_2] \to [R_1]\}_{(u,v)\in E}$, and for every $(u, v) \in E$ the projection $\pi_{v,u}$ is given by:

$$j \in [R_2] : \ \pi_{v,u}(j) = \pi_{v,u}(j + R_2) = {\pi'}_{v,u}(j).$$

In other words, to construct the duplicated instance, we double V's labels set. The new labels added are of the form $j + R_2$ for $j \in [R_2]$, and each new label $j + R_2$ "behaves" like the original label j.

When given a bipartite label cover instance $\mathcal{LC}' = (U, V, E, R, R, \Pi')$ with permutation constraints $\Pi' = \left\{\pi'_{v,u} : [R] \to [R]\right\}_{(u,v)\in E}$, we define the *unique games duplicated-$\mathcal{LC}'$* to be the new bipartite label cover instance $\mathcal{LC} = (U, V, E, 2R, 2R, \Pi)$ with permutation constraints, where $\Pi = \{\pi_{v,u} : [2R] \to [2R]\}_{(u,v)\in E}$, and for every $(u, v) \in E$ the permutation $\pi_{v,u}$ is given by:

$$j \in [R] : \ \pi_{v,u}(j) = {\pi'}_{v,u}(j), \ \pi_{v,u}(j + R) = {\pi'}_{v,u}(j) + R.$$

C. Fourier Analysis

It is well knows that every function $f : \{\pm 1\}^n \to \mathbb{R}$ can be uniquely expressed as a multilinear polynomial (called the *Fourier expansion* of f), that is given by $f(x) = \sum_{S\subseteq[n]} \hat{f}(S)\chi_S(x)$, where for every $S \subseteq [n]$ it holds that $\hat{f}(S) \in \mathbb{R}$, and χ_S is the function $\chi_S : \{\pm 1\}^n \to \{\pm 1\}$ given by $\chi_S(x) = \prod_{i\in S} x_i$.

1) Influences: Let $f : \{\pm 1\}^n \to \mathbb{R}$ be a function, and let $i \in [n]$ be a coordinate. The *influence* of coordinate i on the function f is $Inf_i(f) = \sum_{S:\, i\in S} \hat{f}^2(S)$. Let $d \in \mathbb{N}$. The *d-low-degree influence* of coordinate i on f is $Inf_i^{\leq d}(f) = \sum_{\substack{S:\, i\in S \\ |S|\leq d}} \hat{f}^2(S)$. The influence $Inf_i(f)$ measures how much the function f depends on its i^{th} variable, while the low-degree influence $Inf_i^{\leq d}(f)$ measures this for the low degree part of f. An important property of low-degree influences is that the number of coordinates with a large low-degree influence must be small. In particular, we have the following claim:

Claim II.3. *Let $d \in \mathbb{N}$, $\tau > 0$, and $f : \{\pm 1\}^n \to [-1, 1]$. It holds that $\left|\left\{i \in [n] \,\middle|\, Inf_i^{\leq d}(f) \geq \tau\right\}\right| \leq \frac{d}{\tau}$.*

D. Correlated Probability Spaces

We say that $\left(\prod_{i\in[t]}\Omega_i,\mu\right)$ is a *finite correlated probability space* if μ is a distribution on the finite product set $\prod_{i\in[t]}\Omega_i$. Of a particular interest to us is the case where the correlated space is defined by a measure that is balanced and pairwise independent.

Definition II.4 (Balanceness, Pairwise Independence). Let $\left(\prod_{i\in[t]}\Omega_i,\mu\right)$ be a finite correlated probability space.

We say that μ is *balanced* if, for any $i\in[t]$ and $\omega\in\Omega_i$, it holds that $\Pr_{w\sim\mu}[w_i=\omega]=\frac{1}{|\Omega_i|}$.

We say that μ is *pairwise independent* if, for any $i\neq i'\in[t]$ and $\omega\in\Omega_i,\omega'\in\Omega_{i'}$, it holds that $\Pr_{w\sim\mu}[w_i=\omega\wedge w_{i'}=\omega']=\Pr_{w\sim\mu}[w_i=\omega]\cdot\Pr_{w\sim\mu}[w_{i'}=\omega']$.

III. Covering Hardness of 4LIN

A. PCP Verifier (Proof of Theorem 1)

As usual, we prove Theorem 1 by reduction from label cover. Specifically, we assume to be given a bipartite label cover instance $\mathcal{LC}'=\mathcal{LC}'_{\epsilon,r}$ constructed from a 3SAT formula by the transformation described in Lemma II.2, and construct a PCP verifier that checks proofs for $\mathcal{LC}'$ by only performing 4LIN tests.

Let $\mathcal{LC}'=(U,V,E,R_1,R_2,\Pi')$, $\Pi'=\left\{\pi'_{v,u}:[R_2]\to[R_1]\right\}_{(u,v)\in E}$, be the given instance, and let $\mathcal{LC}=(U,V,E,R_1,2R_2,\Pi)$, $\Pi=\left\{\pi_{v,u}:[2R_2]\to[R_1]\right\}_{(u,v)\in E}$, be the duplicated-$\mathcal{LC}'$ instance (see Section II-B3). A *proof* P for $\mathcal{LC}'$ consists of a collection of truth tables of boolean functions, one for each vertex $v\in V$. Formally, $P=(f_v)_{v\in V}$ where $f_v:\{\pm1\}^{2R_2}\to\{\pm1\}$. The function f_v is, supposedly, the long code encoding of the label assigned to v by a satisfying assignment for $\mathcal{LC}$.

Our verifier's algorithm for checking the proof P is found in Figure 1. The distributions $\mathcal{H}_{\epsilon,u,v,v'}$ on $\left(\{\pm1\}^{2R_2}\right)^4$ used by the verifier are specified in Section III-B2.

Algorithm 1 Ver_ϵ^P

- Randomly select an edge $(u,v)\in_R E$ and a neighbor $v'\in_R\Gamma(u)\subseteq V$.
- Generate a tuple $(x,y,z,w)\in\left(\{\pm1\}^{2R_2}\right)^4$ from the distribution $\mathcal{H}_{\epsilon,u,v,v'}$.
- Accept iff $f(x)f(y)g(z)g(w)=-1$,
  where f and g are the functions in P associated with vertices v and v' (respectively).

Let $\epsilon\in\left(0,\frac{1}{2}\right)$, $k\in\mathbb{N}$, and let $\mathcal{P}=\{P_1,...,P_k\}$ be a set of any k proofs. Define $Rej\left(\mathsf{Ver}_\epsilon^{\mathcal{P}}\right)$ to be the indicator random variable for the rejection of the set of proofs $\mathcal{P}$ by Ver_ϵ. That is, $Rej\left(\mathsf{Ver}_\epsilon^{\mathcal{P}}\right)$ is 1 if none of the proofs in $\mathcal{P}$ satisfies the test selected by Ver_ϵ, and 0 if $\mathcal{P}$ contains a proof that satisfies the test.

We show that the verifier satisfies the following completeness and soundness conditions:

Lemma III.1. Ver *satisfies the following properties:*

- ***Completeness:*** *Let $\epsilon\in\left(0,\frac{1}{2}\right)$ and $r\in\mathbb{N}$. If $\mathsf{OPT}\left(\mathcal{LC}'_{\epsilon,r}\right)=1$, then there exist two proofs, P and Q, such that*
$$\Pr\left[Rej\left(\mathsf{Ver}_\epsilon^{\{P,Q\}}\right)\right]=0.$$
That is, if there is a satisfying assignment for $\mathcal{LC}'_{\epsilon,r}$, then there are 2 proofs that together cover all the tests performed by Ver_ϵ.
Furthermore, each of the proofs P and Q is accepted by Ver_ϵ with probability $1-\epsilon$.
- ***Soundness:*** *For any sufficiently small $\epsilon\in\left(0,\frac{1}{2}\right)$ and sufficiently large $r\in\mathbb{N}$, there exist constants $\delta>0$ and $\xi>0$ that only depend on ϵ (e.g., $\delta=20\sqrt{\epsilon}$ and $\xi=\epsilon^{14}$), such that for any $k\in\mathbb{N}$, the following holds: If there exists a set of proofs $\mathcal{P}$ of size at most k such that*
$$\Pr\left[Rej\left(\mathsf{Ver}_\epsilon^{\mathcal{P}}\right)\right]<\frac{1}{2^k}-\delta.$$
Then $\mathsf{OPT}\left(\mathcal{LC}'_{\epsilon,r}\right)>\xi$.
In particular, if $\mathsf{OPT}\left(\mathcal{LC}'_{\epsilon,r}\right)\le\xi$, then there is no constant number of proofs that together cover all the tests performed by Ver_ϵ.

Note that the soundness property of Lemma III.1 is tight in the sense that k random proofs are expected to cover all but $\frac{1}{2^k}$ fraction of the tests performed by the verifier. We show that no k proofs can do significantly better than k random proofs.

The proof of the completeness part of Lemma III.1 can be found in Section III-C, the proof of the soundness part is omitted. Theorem 1 follows easily from the last lemma.

B. Distributions

Consider the duplicated label cover instance $\mathcal{LC}$. Fix vertices $u\in U$, $v,v'\in\Gamma(u)\subseteq V$, and let $i\in[R_1]$. Let $\mathcal{X}^i,\mathcal{Y}^i=\{\pm1\}^{\pi_{v,u}^{-1}(i)}$ and let $\mathcal{Z}^i,\mathcal{W}^i=\{\pm1\}^{\pi_{v',u}^{-1}(i)}$. For every $i\in[R_1]$ we will have a distribution $\mathcal{H}^i_{\epsilon,u,v,v'}$ on $\Omega^i_{u,v,v'}=\mathcal{X}^i\times\mathcal{Y}^i\times\mathcal{Z}^i\times\mathcal{W}^i$. We think of this space as a correlated space in the sense of Mossel [14], written $\left(\Omega^i_{u,v,v'};\mathcal{H}^i_{\epsilon,u,v,v'}\right)$.

We define $\mathcal{H}_{\epsilon,u,v,v'}$ to be the product distribution $\mathcal{H}_{\epsilon,u,v,v'}=\bigotimes_{i=1}^{R_1}\mathcal{H}^i_{\epsilon,u,v,v'}$ over the domain
$$\Omega_{u,v,v'}=\prod_{i=1}^{R_1}\left(\mathcal{X}^i\times\mathcal{Y}^i\times\mathcal{Z}^i\times\mathcal{W}^i\right)\cong\left(\prod_{i=1}^{R_1}\mathcal{X}^i\right)\times\left(\prod_{i=1}^{R_1}\mathcal{Y}^i\right)\times\left(\prod_{i=1}^{R_1}\mathcal{Z}^i\right)\times\left(\prod_{i=1}^{R_1}\mathcal{W}^i\right).$$

Again, we think of this space as a correlated space $(\Omega_{u,v,v'}; \mathcal{H}_{\epsilon,u,v,v'})$.

1) Our Noise Distribution $\mathcal{N}$: In order to define the distributions $\mathcal{H}^i_{\epsilon,u,v,v'}$, we use the following noise distribution. The distribution $\mathcal{N}_\epsilon(D)$ generates a $2D$-bits string x, such that every coordinate is 1 (noisy) with probability ϵ, but for every $j \in [D]$ it is never the case that both x_j and x_{j+D} are 1.

Definition III.2. Let $\epsilon \in \left[0, \frac{1}{2}\right]$ and $D \in \mathbb{N}$. The distribution $\mathcal{N}_\epsilon(D)$ generates $x = (x_1, \ldots, x_{2D}) \in \{\pm 1\}^{2D}$ as follows: For every $j \in [D]$ independently,

- With probability $1 - 2\epsilon$ set $x_j = x_{j+D} = -1$.
- With probability ϵ set $x_j = -1$ and $x_{j+D} = 1$.
- With probability ϵ set $x_j = 1$ and $x_{j+D} = -1$.

2) The Verifier's Distribution $\mathcal{H}$: Next we define the distribution $\mathcal{H}_{\epsilon,u,v,v'}$ to be used by the verifier: For $D \in \mathbb{N}$ and $a \in \{\pm 1\}$ we denote $a^D = a, \ldots, a$ (the concatenation of a with itself D times). When given $D_1, D_2 \in \mathbb{N}$, we denote $\mathcal{X} = \mathcal{Y} = \{\pm 1\}^{2D_1}$ and $\mathcal{Z} = \mathcal{W} = \{\pm 1\}^{2D_2}$.

Definition III.3. Let $\epsilon \in \left[0, \frac{1}{2}\right]$ and $D_1, D_2 \in \mathbb{N}$. The distribution $\mathcal{H}_\epsilon(D_1, D_2)$ generates $(x_1, ..., x_{2D_1}, y_1, \ldots, y_{2D_1}, z_1, \ldots, z_{2D_2}, w_1, \ldots, w_{2D_2}) \in \mathcal{X} \times \mathcal{Y} \times \mathcal{Z} \times \mathcal{W}$ as follows:

- Select the bits $x_1, ..., x_{2D_1}, z_1, ..., z_{2D_2}$, as well as the auxiliary bit a, independently and uniformly at random.
- Select the auxiliary bits $y'_1, \ldots, y'_{2D_1}$ according to the distribution $\mathcal{N}_\epsilon(D_1)$.
  Select the auxiliary bits $w'_1, \ldots, w'_{2D_2}$ according to the distribution $\mathcal{N}_\epsilon(D_2)$.
  $(x_1, \ldots, x_{2D_1}, z_1, \ldots, z_{2D_2}, a, y'_1, \ldots, y'_{D_1}, w'_1, \ldots, w'_{D_2}$ are all independent).
- Set $y = -x\left(a^{2D_1} \wedge y'\right)$ and $w = -z\left(\left(-a^{2D_2}\right) \wedge w'\right)$. That is, for $j \in [2D_1]$ set $y_j = -x_j\left(a \wedge y'_j\right)$, and for $j \in [2D_2]$ set $w_j = -z_j\left((-a) \wedge w'_j\right)$.

For $i \in [R_1]$ we define $\mathcal{H}^i_{\epsilon,u,v,v'} = \mathcal{H}_\epsilon(d_{i,u,v}, d_{i,u,v'})$ where $d_{i,u,v} = \left|\pi^{-1}_{v,u}(i)\right|$ and $d_{i,u,v'} = \left|\pi^{-1}_{v',u}(i)\right|$. Observe that $\mathcal{H}^i_{\epsilon,u,v,v'}$ can be thought of as simply a distribution on $\left(\{\pm 1\}^{d_{i,v,u}}\right)^2 \times \left(\{\pm 1\}^{d_{i,u,v'}}\right)^2$. As mentioned above, the verifier's distribution is $\mathcal{H}_{\epsilon,u,v,v'} = \bigoplus_{i=1}^{R_1} \mathcal{H}^i_{\epsilon,u,v,v'}$.

C. 4LIN *Completeness*

In this section we prove the completeness property of Lemma III.1. That is, we show that if there exists a satisfying assignment for $\mathcal{LC}'$, then there exist two proofs that together cover the tests of Ver.

Proof of Lemma III.1 (Completeness) Let L be a satisfying assignment for $\mathcal{LC}'$. We construct the two proofs $P = \{f_1^v\}_{v \in V}$ and $Q = \{f_2^v\}_{v \in V}$ for Ver using the assignments L and $L + R_2$ (respectively). That is, $f_1^v, f_2^v : \{\pm 1\}^{2R_2} \to \{\pm 1\}$ satisfy $f_1^v = \chi_{L(v)}$ and $f_2^v = \chi_{L(v)+R_2}$.

Assume that the verifier selects the vertices $u \in U$ and $v, v' \in V$. Denote $i = L(u)$, $j = L(v)$ and $j' = L(v')$. Recall that since L is a satisfying assignment it holds that $\pi_{v,u}(j) = \pi_{v',u}(j') = i$. Let (x, y, z, w) be a possible draw from the distribution $\mathcal{H}_{\epsilon,u,v,v'}$, drawn with $\mathcal{H}^i_{\epsilon,u,v,v'}$ using the auxiliary bit a_i. Note that the tuple (x, y, z, w) induces a test "$f(x) f(y) g(z) g(w) = -1$". For $\ell \in \{1, 2\}$, denote $f_\ell = f_\ell^v$ and $g_\ell = f_\ell^{v'}$. Observe that:

- **If** $a_i = 1$**:** Then $f_\ell(x) \neq f_\ell(y)$ (for example, $f_1(x) = x_j \neq -x_j\left(a_i \wedge y'_j\right) = y_j = f_1(y)$). If additionally $w'_{j'} = -1$ then $g_1(z) = g_1(w)$, and if additionally $w'_{j'+R_2} = -1$ then $g_2(z) = g_2(w)$.
- **If** $a_i = -1$**:** Then $g_\ell(z) \neq g_\ell(w)$ (for example, $g_1(z) = z_{j'} \neq -z_{j'}\left((-a_i) \wedge w'_{j'}\right) = w_{j'} = g_1(w)$). If additionally $y'_j = -1$ then $f_1(x) = f_1(y)$, and if additionally $y'_{j+R_2} = -1$ then $f_2(x) = f_2(y)$.

The above implies that P satisfies the test (x, y, z, w), unless ($a_i = 1$ and $w'_{j'} = 1$) or ($a_i = -1$ and $y'_j = 1$). Similarly, Q satisfies the test (x, y, z, w), unless ($a_i = 1$ and $w'_{j'+R_2} = 1$) or ($a_i = -1$ and $y'_{j+R_2} = 1$). We conclude that each of the proofs P and Q is accepted with probability $1 - \epsilon$.

We now show that at least one of the proofs P and Q satisfies the test (x, y, z, w). We assume that P does not satisfy the test and show that Q does. Assume without loss of generality that $a_i = 1$. Since P does not satisfy the test it must hold that $w'_{j'} = 1$. But since w' is selected according to the distribution $\mathcal{N}_\epsilon(R_2)$, it cannot be the case that both $w'_{j'}$ and $w'_{j'+R_2}$ are 1, thus $w'_{j'+R_2} = -1$. Since $a_i = 1$, this implies $g_2(z) = g_2(w)$ and $f_2(x) \neq f_2(y)$. We conclude that Q satisfies the test (x, y, z, w) and the assertion follows. ■

IV. Covering Dictatorship Test for General Predicates

In this section we develop a general framework for covering dictatorship tests using a given predicate φ, for a large subset of the predicates $\varphi \notin \mathcal{O}$. In particular, in Section IV-A we prove Theorem 2. In Section IV-B we prove a more general version of Theorem 2 (see Lemma IV.6), offering a dictatorship test in the multi-function setting. The general version is used in Section V to obtain a conditional characterization of covering hard predicates.

A. Single Function Dictatorship Test

In this subsection we assume that φ and η are as in Theorem 2, and construct a φ-based covering dictatorship test DICT1, using the distribution η. The test assumes to have an oracle access to the function $f : \{\pm 1\}^{2n} \to \{\pm 1\}$ being tested.

Roughly speaking, we show that our dictatorship test satisfies the following properties: There exist two dictatorships f and g that together cover all the tests made by DICT1. However, any constant number of functions whose every product is "far" from a dictatorship do not cover all the tests made by DICT1. In other words, if a constant number of functions cover all the tests, then there is a subset of these functions whose product is "close" to a dictatorship.

1) The Test: Our dictatorship test uses the distribution $\mathcal{D}_{\epsilon,\eta}$ that is specified next. Let us first define the following noisy version of η:

Definition IV.1. Let $\epsilon \in [0,1]$, let η be a distribution over $\{\pm1\}^t$, and let $\mathcal{U}$ be the uniform distribution over $\{\pm1\}^t$. Define the distribution η'_ϵ generating $y \in \{\pm1\}^t$ as follows: $\forall y : \ \eta'_\epsilon(y) = (1-\epsilon)\,\eta(y) + \epsilon \cdot \mathcal{U}(y)$. That is, in order to draw a t-bits string y from η'_ϵ: With probability $1-\epsilon$ draw y from η, and with probability ϵ draw a random y.

Definition IV.2. Let $\epsilon \in [0,1]$, and let η be a distribution over $\{\pm1\}^t$. Define the distribution $\mathcal{D}_{\epsilon,\eta}$ generating $w = (y,z) \in \left(\{\pm1\}^t\right)^2$ as follows:

$$\forall (y,z) : \ \mathcal{D}_{\epsilon,\eta}(y,z) = \frac{1}{2}\eta(y)\,\eta'_\epsilon(z) + \frac{1}{2}\eta'_\epsilon(y)\,\eta(z).$$

That is, in order to draw a pair (y,z) of t-bits strings from $\mathcal{D}_{\epsilon,\eta}$: With probability $\frac{1}{2}$ draw y from η and z from η'_ϵ, and with probability $\frac{1}{2}$ draw y from η'_ϵ and z from η.

Our dictatorship test is found in Figure 2. For a string s_i, we use the notation $s_{i,j}$ to indicate the j^{th} coordinate of s_i.

Algorithm 2 $\mathsf{DICT1}_\epsilon^f$

- Select $w_1 = (y_1,z_1), \ldots, w_n = (y_n,z_n) \in \left(\{\pm1\}^t\right)^2$ according to the distribution $\mathcal{D}_{\epsilon,\eta}$.
- For $i \in [t]$, let $x_i = y_{1,i}, \ldots, y_{n,i}, z_{1,i}, \ldots, z_{n,i} \in \{\pm1\}^{2n}$.
- Accept iff $\varphi(f(x_1), \ldots, f(x_t)) = -1$.

It will be convenient for us to view the dictatorship test in a matrix notation. For a matrix M, we denote by M_i the i^{th} row of M, and by M^j the j^{th} column of M. Consider the following $2n \times t$ matrix M: The first n rows of the matrix are $y_1, \ldots, y_n$, that is $M_1 = y_1, \ldots, M_n = y_n$. The following n rows are $z_1, \ldots, z_n$, that is $M_{n+1} = z_1, \ldots, M_{2n} = z_n$. Note that the t columns of the obtained matrix are $x_1, \ldots, x_t$, that is $M^1 = x_1, \ldots, M^t = x_t$.

2) Definitions: In order to analyze the dictatorship test we need the following definitions: Let $\epsilon \in (0,1)$, $k \in \mathbb{N}$ and let $\mathcal{F} = \{f_1, \ldots, f_k\}$ be a set of functions $f_\ell : \{\pm1\}^{2n} \to \{\pm1\}$.

We denote by $Rej\left(\mathsf{DICT1}_\epsilon^{\mathcal{F}}\right)$ the indicator random variable for the rejection of the set $\mathcal{P}$ by $\mathsf{DICT1}_\epsilon$. That is, $Rej\left(\mathsf{DICT1}_\epsilon^{\mathcal{F}}\right)$ is 1 if none of the functions f_ℓ in $\mathcal{F}$ passes the test selected by DICT1, and 0 if there exists a function f_ℓ in $\mathcal{F}$ that passes the test.

For $K \subseteq [k]$, we define the product function $f_K : \{\pm1\}^{2n} \to \{\pm1\}$ by $f_K = \prod_{\ell \in K} f_\ell$. The function f_ϕ is the all 1's functions, i.e., for every $x \in \{\pm1\}^{2n}$ it holds that $f_\phi(x) = 1$.

We say that a function is *regular* if none of its coordinates has high influence:

Definition IV.3 (Regularity). Let $d \in \mathbb{N}$, $\tau \in [0,1]$, and let $f : \{\pm1\}^n \to \{\pm1\}$ be a function. We say that f is (d,τ)*-regular*, if $\max_{j\in[n]}\left\{Inf_j^{\le d}(f)\right\} \le \tau$. Let $\mathcal{F} = \{f_1, \ldots, f_k\}$ be a set of functions $f_\ell : \{\pm1\}^n \to \{\pm1\}$. We say that $\mathcal{F}$ is (d,τ)*-regular* if for every subset $K \subseteq [k]$ it holds that f_K is (d,τ)-regular.

3) Test Analysis: We are now ready to state our result regarding DICT1. The following lemma clearly implies Theorem 2.

Lemma IV.4. *Let $\varphi : \{\pm1\}^t \to \{\pm1\}$ be a predicate satisfying $\varphi \notin \mathcal{O}$. Assume that there exists a balanced, pairwise independent distribution η on the support of φ. Then,* DICT1 *satisfies the following properties:*

- ***Completeness:*** *For any $\epsilon \in (0,1)$ the following holds. Let $j \in [n]$ and let $f, g : \{\pm1\}^{2n} \to \{\pm1\}$ be the two dictatorships $f = \chi_j$ and $g = \chi_{j+n}$. Then*

  $$\Pr\left[Rej\left(\mathsf{DICT1}_\epsilon^{\{f,g\}}\right)\right] = 0.$$

  In particular, there exist two dictatorships that cover all the tests performed by DICT1*.*
  Furthermore, each of the functions f and g is accepted by $\mathsf{DICT1}_\epsilon$ with probability $1 - \frac{\epsilon}{2}$.
- ***Soundness:*** *For any $\epsilon \in (0,1)$ and $k \in \mathbb{N}$, there exist constants $d \in \mathbb{N}$ and $\tau > 0$ that only depend on ϵ, t and k, such that the following holds. Let $\mathcal{F} = \{f_1, \ldots, f_k\}$ be a set of functions $f_\ell : \{\pm1\}^{2n} \to \{\pm1\}$. Assume that $\mathcal{F}$ is (d,τ)-regular. Then*

  $$\Pr\left[Rej\left(\mathsf{DICT1}_\epsilon^{\mathcal{F}}\right)\right] > \frac{1}{2^{10kt}}.$$

  In particular, if a set of a constant number of functions covers all the tests performed by DICT1*, then the set is not (d,τ)-regular.*

B. General Dictatorship Test

In this section we offer a more general covering dictatorship test. Instead of getting oracle access to a single function $f : \{\pm1\}^{2n} \to \{\pm1\}$, the general dictatorship test gets access to a family of functions $F = \left\{f^1, \ldots, f^r\right\}$, $r \ge t \in \mathbb{N}$, $f : \{\pm1\}^{2n} \to \{\pm1\}$. Intuitively, the test aims at checking whether all the functions $f^1, \ldots, f^r$ are the same dictatorship.

The general dictatorship test is used in Section V. Roughly speaking, the $f^1, \ldots, f^r$ functions considered by Section V are the long code encodings of the $r \gg t$ neighbors $v_1, \ldots, v_r$ of a single vertex u in a label cover instance.

1) The Test: Our general dictatorship test DICT is found in Figure 3. It is easy to see that if F contains r copies of the same function f, i.e., $f^1 = \cdots = f^r = f$, then DICT_ϵ^F operates the same as $\mathsf{DICT1}_\epsilon^f$.

Algorithm 3 $\mathsf{DICT}_\epsilon^{F=\left\{f^1,\ldots,f^r\right\}}$

- Select $w_1 = (y_1, z_1), \ldots, w_n = (y_n, z_n) \in \left(\{\pm 1\}^t\right)^2$ according to the distribution $\mathcal{D}_{\epsilon,\eta}$.
- For $i \in [t]$, let $x_i = y_{1,i}, \ldots, y_{n,i}, z_{1,i}, \ldots, z_{n,i} \in \{\pm 1\}^{2n}$.
- Select a random set of t *different* indices $I = \{i_1, \ldots, i_t\} \subseteq [r]$ (selection with no repetitions).
- Accept iff $\varphi\left(f^{i_1}(x_1), \ldots, f^{i_t}(x_t)\right) = -1$.

---

2) Definitions: In order to analyze the dictatorship test we need the following definitions: Let $\epsilon \in (0,1)$, $k, r \in \mathbb{N}$, and let $\mathcal{F}$ be a set of sets of functions $\mathcal{F} = \{F_1, \ldots, F_k\}$, $F_\ell = \left\{f_\ell^1, \ldots, f_\ell^r\right\}$, $f_\ell^i : \{\pm 1\}^{2n} \to \{\pm 1\}$.

We denote by $Rej\left(\mathsf{DICT}_\epsilon^{\mathcal{F}}\right)$ the indicator random variable for the rejection of the set $\mathcal{F}$ by DICT. That is, $Rej\left(\mathsf{DICT}_\epsilon^{\mathcal{F}}\right)$ is 1 if none of the sets F_ℓ in $\mathcal{F}$ passes the test selected by DICT, and 0 if there exists a set F_ℓ in $\mathcal{F}$ that passes the test.

We define the *cross influence* of a pair of function as follows:

Definition IV.5 (Cross Influence). Let $d \in \mathbb{N}$, and let $f, g : \{\pm 1\}^n \to \{\pm 1\}$ be a pair of function. We denote by $XInf_d(f,g)$ the *d-cross influence* of f and g:

$$XInf_d(f,g) = \max_{j \in [n]} \left\{ \min \left\{ Inf_j^{\leq d}(g), Inf_j^{\leq d}(f) \right\} \right\}.$$

For $i \in [r]$ and $K \subseteq [k]$ we define the product function $f_K^i : \{\pm 1\}^{2n} \to \{\pm 1\}$ by $f_K^i = \prod_{\ell \in K} f_K^i$. The function f_ϕ^i is the all 1's functions, i.e., for every $x \in \{\pm 1\}^{2n}$ it holds that $f_\phi^i(x) = 1$.

Let $d \in \mathbb{N}$ and $\tau \in [0,1]$. Let $(i, i') \in [r]^2$, $i \neq i'$, be a pair of indices. We say that (i,i') is $(d,\tau,\mathcal{F})$*-cross regular*, if for every two sets $K, K' \subseteq [k]$ it holds that $XInf_d\left(f_K^i, f_{K'}^{i'}\right) \leq \tau$. Let $I \subseteq [r]$ be a set of indices. We say I is $(d,\tau,\mathcal{F})$*-cross regular* if every pair $(i,i') \in I^2$, $i \neq i'$, is $(d,\tau,\mathcal{F})$-cross regular.

3) Test Analysis: We are now ready to state our result regarding DICT:

Lemma IV.6. *Let $\varphi : \{\pm 1\}^t \to \{\pm 1\}$ be a predicate satisfying $\varphi \notin \mathcal{O}$. Assume that there exists a balanced, pairwise independent distribution η on the support of φ. Then,* DICT *satisfies the following properties:*

- ***Completeness:*** *For any $\epsilon \in (0,1)$ the following holds. Let $j \in [n]$ and let $f^i, g^i : \{\pm 1\}^{2n} \to \{\pm 1\}$, $i \in [r]$, be the following functions $f^1 = \cdots = f^r = \chi_j$, $g^1 = \cdots = g^r = \chi_{j+n}$. Let $F = \left\{f^1, \ldots, f^r\right\}$ and $G = \left\{g^1, \ldots, g^r\right\}$. Then*

$$\Pr\left[Rej\left(\mathsf{DICT}_\epsilon^{\{F,G\}}\right)\right] = 0.$$

Furthermore, each of the sets F and G is accepted by DICT_ϵ with probability $1 - \frac{\epsilon}{2}$.
- ***Soundness:*** *For any $\epsilon \in (0,1)$ and $k \in \mathbb{N}$, there exist $r, d \in \mathbb{N}$ and $\tau > 0$ that only depend on ϵ, t, and k, such that the following holds. Let $\mathcal{F}$ be a set of sets of functions $\mathcal{F} = \{F_1, \ldots, F_k\}$, $F_\ell = \left\{f_\ell^1, \ldots, f_\ell^r\right\}$, $f_\ell^i : \{\pm 1\}^{2n} \to \{\pm 1\}$. Let $\alpha = \frac{1}{2^{10kt}}$, and assume that at least $1-\alpha$ fraction of the t-elements sets $I \subseteq [r]$* are *$(d,\tau,\mathcal{F})$-cross regular. Then*

$$\Pr\left[Rej\left(\mathsf{DICT}_\epsilon^{\mathcal{F}}\right)\right] > \frac{1}{2^{10kt}}.$$

It is easy to see that Lemma IV.4 is a special case of Lemma IV.6. Thus, we will only prove Lemma IV.6. The proof of the completeness part of Lemma IV.6 can be found in Section IV-C, the proof of the soundness part is omitted. For the rest of the text we fix t, ϵ, φ and η. We omit the ϵ and η sub-indices and write $\mathsf{DICT} = \mathsf{DICT}_\epsilon$, $\eta' = \eta'_\epsilon$ and $\mathcal{D} = \mathcal{D}_{\epsilon,\eta}$.

C. Completeness

Proof of Lemma IV.6 (Completeness) Recall the matrix M defined at the end of Subsection IV-A1, after the algorithm DICT1. For $i \in [t]$, it holds that x_i is the i^{th} column of M. For $j \in [n]$ it holds that y_j is the j^{th} row of M, and z_j is row number $j+n$ of M. Also recall that there exists $j \in [n]$ such that for every $i \in [r]$ it holds that $f^i = \chi_j$ and $g^i = \chi_{j+n}$.

When running DICT^F we compare the following value to -1: $\varphi\left(f^{i_1}(x_1), \ldots, f^{i_t}(x_t)\right) = \varphi\left(\chi_j(x_1), \ldots, \chi_j(x_t)\right) = \varphi\left(M_{j,1}, \ldots, M_{j,t}\right) = \varphi\left(M_j\right) = \varphi\left(y_j\right)$. When running DICT^G we compare the following value to -1: $\varphi\left(g^{i_1}(x_1), \ldots, g^{i_t}(x_t)\right) = \varphi\left(\chi_{j+n}(x_1), \ldots, \chi_{j+n}(x_t)\right) = \varphi\left(M_{j+n,1}, \ldots, M_{j+n,t}\right) = \varphi\left(M_{j+n}\right) = \varphi\left(z_j\right)$

Recall that (y_j, z_j) was drawn from $\mathcal{D}$. Thus, either y_j or z_j was drawn from η, implying that at least one of them is in support of φ. Hence, either $\varphi(y_j) = -1$ or $\varphi(z_j) = -1$, and at least one of F and G is accepted by DICT.

Moreover, observe that F is only rejected by DICT if y_j is not in the support of φ. This can only happen if $\mathcal{D}$ samples y_i using η', and η' samples y_i using the uniform distribution (instead of using η). The probability of this event is $\frac{\epsilon}{2}$. Therefore, F (and similarly also G) is accepted with probability at least $1 - \frac{\epsilon}{2}$. ■

V. Characterization of Covering-Hard Predicates

In this section we prove covering unique games hardness for a large subset of the predicates $\varphi \notin \mathcal{O}$. Formally, we prove Theorem 4 under Conjecture 3 (covering UGC).

A. Discussion of our Covering Unique Games Conjecture

We would like to follow the lines of [2] and get a conditional hardness result using our dictatorship test. A natural attempt at formulating a covering conjecture would be to require in the yes case the existence of c assignments that together cover all the edges of the given label cover instance, where c is some absolute constant. Unfortunately, we were only able to derive a hardness result using a stronger version of the conjecture. Specifically, in the yes case we require the existence of c assignments such that for every vertex $u \in U$, at least one of the assignments satisfies all the edges touching u. We mention that Khot and Regev [12] show that a similar conjecture in the max-CSP setting is equivalent to the unique games conjecture.

Our conjecture is clearly false with $c = 1$, but as far as we know may be true with even $c = 2$. The conjecture is incomparable to the unique games conjecture (our completeness does not require any single assignment to satisfy a large fraction of edges). However it clearly implies the unique games conjecture with completeness $\frac{1}{c}$ (instead of $1 - \epsilon$).

B. PCP Verifier (Proof of Theorem 4)

As usual, we prove Theorem 4 by reduction from the covering unique games conjecture (Conjecture 3). Specifically, we assume to be given a bipartite label cover instance $\mathcal{LC}'$ with permutation constraints, and construct a PCP verifier that checks proofs for $\mathcal{LC}'$ by only performing φ-tests.

Let $\mathcal{LC}' = (U, V, E, R, R, \Pi')$, $\Pi' = \left\{\pi'_{v,u} : [R] \to [R]\right\}_{(u,v)\in E}$, be the given instance, and let $\mathcal{LC} = (U, V, E, 2R, 2R, \Pi)$, $\Pi = \left\{\pi_{v,u} : [2R] \to [2R]\right\}_{(u,v)\in E}$, be the unique games duplicated-$\mathcal{LC}'$ instance (see Section II-B3). A *proof* P for $\mathcal{LC}'$ consists of a collection of truth tables of boolean functions, one for each vertex $v \in V$. Formally, $P = (f_v)_{v\in V}$ where $f_v : \{\pm1\}^{2R} \to \{\pm1\}$. The function f_v is, supposedly, the long code encoding of the label assigned to v by a satisfying assignment for $\mathcal{LC}$.

Our verifier's algorithm for checking the proof P is found in Figure 4. The algorithm uses the following definition. For a function $f : \{\pm1\}^{2R} \to \{\pm1\}$ and a permutation $\pi : [2R] \to [2R]$ we define the function $f\pi : \{\pm1\}^{2R} \to \{\pm1\}$ by $f\pi(x) = f\left(x_{\pi(1)}, \ldots, x_{\pi(2R)}\right)$.

Algorithm 4 Ver_ϵ^P

- Randomly select a vertex $u \in_R U$.
- Run DICT_ϵ^F for $F = \{f^v\pi_{v,u}\}_{v\in\Gamma(u)}$, where f^v is the function in P associated with vertex v.

As before, we define $Rej\left(\mathsf{Ver}_\epsilon^{\mathcal{P}}\right)$ to be the indicator random variable for the rejection of the set of proofs $\mathcal{P} = \{P_1, ..., P_k\}$ by Ver_ϵ. Theorem 4 is an easy corollary of the following lemma:

Lemma V.1. *Let $c \in \mathbb{N}$ be the constant promised by the covering unique games conjecture (Conjecture 3). Let $\varphi : \{\pm1\}^t \to \{\pm1\}$ be a predicate satisfying $\varphi \notin \mathcal{O}$. Assume that there exists a balanced, pairwise independent distribution η on the support of φ. Then,* Ver *satisfies the following properties:*

- ***Completeness:*** *Assume that there exist c assignments such that for every vertex $u \in U$, at least one of the assignments satisfies all the edges touching u. Then, there exists a set $\mathcal{P}$ of at most $2c$ proofs such that*
$$\Pr\left[Rej\left(\mathsf{Ver}_\epsilon^{\mathcal{P}}\right)\right] = 0.$$
In particular, if c assignments cover all the edges of $\mathcal{LC}'$ (in the above sense), then there are $2c$ proofs that together cover all the tests performed by Ver_ϵ.
- ***Soundness:*** *For a sufficiently small $\epsilon > 0$ and $k \in \mathbb{N}$, there exists a constant $\xi > 0$ that only depends on ϵ, t and k, such that the following holds: Assume that there exists a set $\mathcal{P}$ of at most k proofs such that*
$$\Pr\left[Rej\left(\mathsf{Ver}_\epsilon^{\mathcal{P}}\right)\right] < \frac{1}{2 \cdot 2^{10kt}}.$$
Then, $\mathsf{OPT}\left(\mathcal{LC}'\right) > \xi$.
In particular, if $\mathsf{OPT}\left(\mathcal{LC}'\right) \le \xi$, *then there is no constant number of proofs that together cover all the tests performed by* Ver_ϵ.

Proof of Lemma V.1 (Completeness) Let $L_1, \ldots, L_c : U \cup V \to [R]$ be the c promised assignments for $\mathcal{LC}'$. Specifically, for every $u \in U$ there exists an assignment L_ℓ, $\ell \in [c]$, that satisfies all the edges touching u.

For each $\ell \in [c]$, we construct the two proofs $P_\ell = \{f_\ell^v\}_{v\in V}$ and $Q_\ell = \{g_\ell^v\}_{v\in V}$ for Ver using the assignments L_ℓ and $L_\ell + R$ (respectively). That is, $f_\ell^v, g_\ell^v : \{\pm1\}^{2R} \to \{\pm1\}$ satisfy $f_\ell^v = \chi_{L_\ell(v)}$ and $g_\ell^v = \chi_{L_\ell(v)+R}$. We denote $\mathcal{P} = \{P_\ell, Q_\ell\}_{\ell\in[c]}$. We next show that $\mathsf{Ver}^{\mathcal{P}}$ always accepts.

Fix a vertex $u \in U$, and assume that u was selected during the execution of Ver. Let L_ℓ, $\ell \in [c]$, be an assignment that satisfies all the edges touching u. When running $\mathsf{Ver}^{\{P_\ell,Q_\ell\}}$, the verifier runs $\mathsf{DICT}^{\{F_{\ell,u},G_{\ell,u}\}}$ for the sets $F_{\ell,u} = \{f_\ell^v\pi_{v,u}\}_{v\in\Gamma(u)}$ and $G_{\ell,u} = \{g_\ell^v\pi_{v,u}\}_{v\in\Gamma(u)}$. For every $v \in \Gamma(u)$ it holds that $f_\ell^v\pi_{v,u} = \chi_{L_\ell(v)}\pi_{v,u} = \chi_{\pi_{v,u}(L_\ell(v))} = \chi_{L_\ell(u)}$, and $g_\ell^v\pi_{v,u} = \chi_{L_\ell(v)+R}\pi_{v,u} = \chi_{\pi_{v,u}(L_\ell(v)+R)} = \chi_{L_\ell(u)+R}$. Using the completeness property of Lemma IV.6, $\mathsf{DICT}^{\{F_{\ell,u},G_{\ell,u}\}}$ always accepts. This implies that $\mathsf{Ver}^{\{P_\ell,Q_\ell\}}$ accepts whenever u is chosen. Therefore, $\mathsf{Ver}^{\mathcal{P}} = \mathsf{Ver}^{\{P_\ell,Q_\ell\}_{\ell\in[c]}}$ always accepts, and the assertion follows. ∎

Proof of Lemma V.1 (Soundness) Let $k \in \mathbb{N}$ and let $\mathcal{P}$ be a set k proofs $\mathcal{P} = \{P_1, \dots, P_k\}$, $P_\ell = \{f_\ell^v\}_{v \in V}$, $f_\ell^v : \{\pm 1\}^{2R} \to \{\pm 1\}$, for which $\Pr\left[Rej\left(\mathsf{Ver}^{\mathcal{P}}\right)\right] < \frac{1}{2 \cdot 2^{10kt}}$. We wish to show that $\mathsf{OPT}\left(\mathcal{LC}'\right) > \xi$, for some constant ξ that only depends on ϵ, t and k.

For simplicity of exposition we assume that $\mathcal{LC}'$ (and therefore also $\mathcal{LC}$) is U-regular, that is $\forall u, u' \in U$: $|\Gamma(u)| = |\Gamma(u')| = r$ and that r is sufficiently large (as required by soundness property of Lemma IV.6).

Fix a vertex $u \in U$, and let $v_1, \dots, v_r \in \Gamma(u)$ be its r neighbors. For $i \in [r]$ and $\ell \in [k]$, denote $g_\ell^{v_i} = f_\ell^{v_i} \pi_{v_i,u}$. Let $F_\ell^u = \{g_\ell^{v_1}, \dots, g_\ell^{v_r}\}$ and $\mathcal{F}^u = \{F_1^u, \dots, F_k^u\}$. It holds that $\Pr\left[Rej\left(\mathsf{Ver}^{\mathcal{P}}\right)\right] = \mathbb{E}_{u \in U}\left[\Pr\left[Rej\left(\mathsf{DICT}^{\mathcal{F}^u}\right)\right]\right]$.

Since $\Pr\left[Rej\left(\mathsf{Ver}^{\mathcal{P}}\right)\right] < \frac{1}{2 \cdot 2^{10kt}}$ there exists a subset $U' \subseteq U$, $|U'| \geq \frac{1}{2}|U|$, such that for every $u \in U'$ it holds that $\Pr\left[Rej\left(\mathsf{DICT}^{\mathcal{F}^u}\right)\right] \leq \frac{1}{2^{10kt}}$. We call the vertices in U' *good* vertices.

Fix a good vertex $u \in U'$. Using the soundness property of Lemma IV.6, for some $d \in \mathbb{N}$ and $\tau > 0$ (functions of ϵ, t and k), it holds that at least $\alpha = \frac{1}{2^{10kt}}$ fraction of the t-elements sets $I \subseteq [r]$ are *not* $(d, \tau, \mathcal{F}^u)$-cross regular. Meaning that there exists a pair $(i, i') \in I^2$, $i \neq i'$, that is not $(d, \tau, \mathcal{F}^u)$-cross regular. We say that such a pair is *cross influential for I* with respect to $\mathcal{F}^u$. We call a pair *cross influential* with respect to $\mathcal{F}^u$, if it cross influential with respect to $\mathcal{F}^u$ for at least one set I.

Denote by $XInfPairs^u \subseteq [r]^2$ the set of cross influential pairs with respect to $\mathcal{F}^u$. Formally, $XInfPairs^u = \left\{(i, i') \in [r]^2 \,|\, \exists K, K' \subseteq [k] : XInf_d\left(g_K^{v_i}, g_{K'}^{v_{i'}}\right) \geq \tau\right\}$. We claim that the set $XInfPairs^u$ contains at least $\frac{\alpha}{t^2}$ fractions of the pairs in $[r]^2$. The following is a way of choosing a random pair $(i, i') \in [r]^2$: First select a t-elements set $I \subseteq [r]$, then select a random pair $(i, i') \in I^2$. The selected set I has a cross influential pair with probability is at least α. Each of the pairs in I is selected with probability $\frac{1}{t^2}$. Thus, the selected pair is cross influential with probability at least $\frac{\alpha}{t^2}$.

Obtaining a good labeling: We next construct a good labeling for the duplicated label cover instance $\mathcal{LC}$. Since every assignment for $\mathcal{LC}$ naturally induces an assignment for $\mathcal{LC}'$ with at least the same value, the claim of the lemma follows.

Consider the following labeling L for $\mathcal{LC}$: The set of candidate assignments for vertex $v \in V$ is given by

$$C(v) = \left\{j \in [2R] \,|\, \exists K \subseteq [k] : Inf_j^{\leq d}\left(f_K^v\right) \geq \tau\right\}.$$

Note that, using Claim II.3, $|C(v)| \leq \frac{d}{\tau} \cdot 2^k$. Define a labeling L by picking, for each $v \in V$ a label $L(v)$ uniformly at random from $C(v)$ (or an arbitrary label if $C(v)$ is empty). For $u \in U$, randomly select $v' \in_R \Gamma(u)$ and set $L(u)$ to $\pi_{v',u}(L(v'))$.

Let (u, v) be an edge of $\mathcal{LC}$, where $u \in U'$ is good. The probability that the edge (u, v) is satisfied by L is

$$\Pr\left[L(u) = \pi_{v,u}(L(v))\right] = \Pr\left[\pi_{v',u}(L(v')) = \pi_{v,u}(L(v))\right].$$

Recall that with probability at least $\frac{\alpha}{t^2}$ it holds that (v, v') is a cross influential pair with respect to $\mathcal{F}^u$, i.e., $v = v_i$, $v' = v_{i'}$, and $(i, i') \in XInfPairs^u$. Therefore, with probability at least $\frac{\alpha}{t^2}$, it holds that $\pi_{v',u}(C(v')) \cap \pi_{v,u}(C(v)) \neq \phi$. Conclude that the edge (u, v) is satisfied with probability at least $\frac{\alpha}{t^2} \cdot \frac{1}{|C(v)||C(v')|} \geq \frac{\alpha}{t^2}\left(\frac{\tau}{d \cdot 2^k}\right)^2$.

Since we assume that $\mathcal{LC}$ is U-regular, and since $|U'| \geq \frac{1}{2}|U|$, it holds that $\frac{1}{2}$ of the edges (u, v) have $u \in U'$. Therefore, a random edge of $\mathcal{LC}$ is satisfied by L with probability at least $\frac{1}{2} \cdot \frac{\alpha}{t^2} \cdot \left(\frac{\tau}{d \cdot 2^k}\right)^2 \geq \frac{\tau^2}{d^2 \cdot 2^{20kt}} = \xi$. ∎

VI. Hardness of Approximate Coloring and Covering Random CSP Instances

An outstanding open question is to approximate the number of colors required to color a given $O(1)$-colorable graph or hypergraph. While it is known to be hard to color a $O(1)$-colorable hypergraph with a polylogarithmic number of colors, the best known algorithm requires a polynomial number of colors. Thus, there is an exponential gap between the best lower and upper bounds. In the covering language this is almost[1] equivalent to the question of approximating the covering number of an $O(1)$-coverable NAE instance. We next study this question in relation to the hardness of random CSP instances.

In a seminal paper, Feige [5] studies the relation between hardness of random instances of 3SAT and the hardness of approximation problems, including some notorious problems for which neither algorithms nor hardness are known. In that paper he states a hypothesis about no polynomial time algorithm being able to distinguish between a random 3SAT and a satisfiable one. More accurately,

Hypothesis VI.1 (Feige's Hypothesis 1 [5]). *There is no polynomial time algorithm that outputs* `typical` *for most* 3CNF-*CSP instances with n variables and $m = \Delta \cdot n$ clauses, and never outputs* `typical` *on a satisfiable instance; even when Δ is an arbitrarily large constant independent of n.*

We formulate an analogous hypothesis about the hardness of distinguishing between random and 2-coverable 4LIN-CSP instances (Hypothesis 5), and a weaker hypothesis about φ-CSP instances for some predicate φ (Hypothesis VI.2). We prove that both of these hypotheses imply the hardness of approximate coloring of hypergraphs. We show a direct connection between the density Δ in the hypothesis and the inapproximability factor in the result. When our

[1] It is not exactly equivalent since the NAE formulation allows negations of variables whereas the coloring formulation does not.

4LIN hypothesis is pushed to extreme, it implies hardness of approximate coloring to within polynomial factors.

Hypothesis VI.2 (**Covering φ Hypothesis, with density parameter Δ**). *There are some universal constants $t, c \in \mathbb{N}$ and a predicate $\varphi : \{\pm 1\}^t \to \{\pm 1\}$ such that no polynomial time algorithm outputs* `typical` *for most φ-CSP instances with n variables and $m = \Delta \cdot n$ clauses, and never outputs* `typical` *for a c-coverable φ-CSP instance.*

Our main theorem of this section is Theorem VI.3 (generalized restatement of Theorem 6):

Theorem VI.3. *If Hypothesis VI.2 holds with parameters c, t and density Δ such that $c \ll \log \Delta$ then it is hard to distinguish if a given t-uniform hypergraph is 2^c-colorable or $\Delta^{\Omega(1)}$ colorable.*

In particular, Hypothesis 5 with density parameter $\Delta = n^\delta$ for some positive $\delta > 0$ implies that it is hard to decide if a 4-uniform hypergraph is 4-colorable or requires at least a polynomial number of colors.

The main ingredient in the proof of Theorem VI.3 is the following claim that shows that the covering number of a random φ-CSP is proportional to its log-density, as long as $\varphi \notin \mathcal{O}$. (Recall that for any $\varphi \in \mathcal{O}$, the covering number of any φ-CSP is at most 2). The theorem's proof is omitted.

Claim VI.4. *Let $\varphi \notin \mathcal{O}$ and let $\mathcal{C}$ be a random instance of φ-CSP, with n variables and $m = \Delta \cdot n$ constraints. Then $\nu(\mathcal{C}) \geq \Omega(\log \Delta)$, except with probability exponentially small in n.*

Proof Let us first assume that φ is the NAE predicate on t variables and that all occurrences of φ are unsigned, i.e., without negations of variables. Fix an CSP instance $\mathcal{C}$, and let $L_1, \dots, L_k \in \{\pm 1\}^n$ be any set of $k \in \mathbb{N}$ assignments for $\mathcal{C}$. It is not hard to see that there must be a subset $S \subseteq [n]$, such that each assignment L_ℓ, $\ell \in [k]$, is constant on S (either all 1s or all -1s), and such that $|S| \geq n \cdot 2^{-k}$. The reason is that each of the assignments L_ℓ partitions the n variables into two sets: Variables that are assigned the value 1, and variable that are assigned the value -1.

If the given instance $\mathcal{C}$ has a constraint fully contained in S then $L_1, \dots, L_k$ do not cover it. The probability that a randomly chosen constraint is contained in a set of size $n \cdot 2^{-k}$ is 2^{-kt} where t is the arity of the constraint. The probability that out of m constraints of $\mathcal{C}$ none landed inside S is $\left(1 - 2^{-kt}\right)^m \approx \exp\left(\frac{-m}{2^{kt}}\right)$, and if we multiply this by the number 2^{kn} of possibilities to choose k assignments and using a union bound we get

$$\Pr_I\left[\nu(I) \leq k\right] \leq \exp\left(\frac{-m}{2^{kt}}\right) \cdot 2^{kn} = \exp\left(-n \cdot \left(\frac{\Delta}{2^{kt}} - k\right)\right).$$

Clearly if $\Delta > 2^{2kt}$ then $\frac{\Delta}{2^{kt}} - k > 1$ which causes the above probability to be exponentially small. In our case t is fixed, and so this proves that $\nu(\mathcal{C}) \geq \Omega(\log \Delta)$ with high probability.

It remains to justify the assumption that φ is the NAE predicate. This simply follows from the fact that for every $\varphi \notin \mathcal{O}$ there is some signed-NAE predicate that contains it, see Claim II.1. The unsigned assumption means that we've proven that even covering the unsigned part of the instance is already hard, assuming that there are many unsigned constraints. But this is indeed the case as the number of unsigned constraints is expected to be $m \cdot 2^{-t}$. ■

References

[1] S. Arora, E. Chlamtac, and M. Charikar, "New approximation guarantee for chromatic number," in *STOC*, 2006, pp. 215–224.

[2] P. Austrin and E. Mossel, "Approximation resistant predicates from pairwise independence," *Computational Complexity*, vol. 18, no. 2, pp. 249–271, 2009.

[3] A. Blum and D. R. Karger, "An $\tilde{O}\left(n^{3/14}\right)$-coloring algorithm for 3-colorable graphs," *Information Processing Letters*, vol. 61, no. 1, pp. 49–53, 1997.

[4] I. Dinur, O. Regev, and C. D. Smyth, "The hardness of 3-uniform hypergraph coloring," *Combinatorica*, vol. 25, no. 5, pp. 519–535, 2005.

[5] U. Feige, "Relations between average case complexity and approximation complexity," in *STOC*, 2002, pp. 534–543.

[6] U. Feige and E. Ofek, "Easily refutable subformulas of large random 3CNF formulas," *Theory of Computing*, vol. 3, no. 1, pp. 25–43, 2007.

[7] V. Guruswami, J. Håstad, and M. Sudan, "Hardness of approximate hypergraph coloring," *SIAM Journal on Computing*, vol. 31, no. 6, pp. 1663–1686, 2002.

[8] J. Håstad, "Some optimal inapproximability results," *Journal of the ACM*, vol. 48, no. 4, pp. 798–859, 2001.

[9] D. R. Karger, R. Motwani, and M. Sudan, "Approximate graph coloring by semidefinite programming," *Journal of the ACM*, vol. 45, no. 2, pp. 246–265, 1998.

[10] S. Khot, "Hardness results for approximate hypergraph coloring," in *STOC*, 2002, pp. 351–359.

[11] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell, "Optimal inapproximability results for max-cut and other 2-variable CSPs?" *SIAM Journal on Computing*, vol. 37, no. 1, pp. 319–357, 2007.

[12] S. Khot and O. Regev, "Vertex cover might be hard to approximate to within 2-epsilon," *Journal of Computer and System Sciences*, vol. 74, no. 3, pp. 335–349, 2008.

[13] S. Khot and R. Saket, "A 3-query non-adaptive pcp with perfect completeness," in *IEEE Conference on Computational Complexity*, 2006, pp. 159–169.

[14] E. Mossel, "Gaussian bounds for noise correlation of functions and tight analysis of long codes," in *FOCS*, 2008, pp. 156–165.

[15] E. Mossel, R. O'Donnell, and K. Oleszkiewicz, "Noise stability of functions with low influences: invariance and optimality," in *FOCS*, 2005, pp. 21–30.

[16] R. O'Donnell and J. Wright, "A new point of NP-hardness for unique games," in *STOC*, 2012, pp. 289–306.

[17] R. O'Donnell and Y. Wu, "Conditional hardness for satisfiable 3-CSPs," in *STOC*, 2009, pp. 493–502.