Electronic Colloquium on Computational Complexity, Report No. 116 (2016)

Approximating CSPs using LP Relaxation*

Subhash Khot[†] Rishi Saket[‡]

Abstract

This paper studies how well the standard LP relaxation approximates a k-ary constraint satisfaction problem (CSP) on label set $[L]$. We show that, assuming the Unique Games Conjecture, it achieves an approximation within $O(k^3 \cdot \log L)$ of the optimal approximation factor. In particular we prove the following hardness result: let $\mathcal{I}$ be a k-ary CSP on label set $[L]$ with constraints from a *constraint class* $\mathcal{C}$, such that it is a (c, s)-integrality gap for the standard LP relaxation. Then, given an instance $\mathcal{H}$ with constraints from $\mathcal{C}$, it is NP-hard to decide whether,

$$\mathsf{opt}(\mathcal{H}) \geq \Omega\left(\frac{c}{k^3 \log L}\right), \quad \text{or} \quad \mathsf{opt}(\mathcal{H}) \leq 4 \cdot s,$$

assuming the Unique Games Conjecture. We also show the existence of an efficient LP rounding algorithm Round such that a lower bound for it can be translated into a similar (but weaker) hardness result. In particular, if there is an instance from a *permutation invariant* constraint class $\mathcal{C}$ which is a (c, s)-*rounding gap* for Round, then given an instance $\mathcal{H}$ with constraints from $\mathcal{C}$, it is NP-hard to decide whether,

$$\mathsf{opt}(\mathcal{H}) \geq \Omega\left(\frac{c}{k^3 \log L}\right), \quad \text{or} \quad \mathsf{opt}(\mathcal{H}) \leq O\left((\log L)^k\right) \cdot s,$$

assuming the Unique Games Conjecture.

1 Introduction

A k-ary constraint satisfaction problem (CSP) over label set $[L]$ consists of a set of vertices and a set of k-uniform ordered hyperedges. For each hyperedge there is a constraint specifying the k-tuples of labels to the vertices in it that satisfy the hyperedge. The goal is to efficiently compute an assignment that satisfies the maximum number of hyperedges. This general definition includes many problems studied in computer science and combinatorial optimization such as MAXIMUM CUT, MAX-k-SAT and MAX-k-LIN$[q]$. Investigating the approximability of these problems has motivated a significant body of research.

One of the well studied methods of approximating a CSP is via the Linear Programming (LP) relaxation of the corresponding integer program[1]. For example, in its most basic formulation the LP relaxation gives a 2-approximation for MAXIMUM CUT and can do no better. On the other hand the seminal work of Goemans and Williamson [6] gave a 1.13823-approximation for MAXIMUM CUT using a semi-definite

*An extended abstract of this paper appeared in the Proceedings of The 42nd International Colloquium on Automata, Languages, and Programming (ICALP 2015).

[†]Computer Science Department, New York University, USA. Email: khot@cims.nyu.edu. Research supported by NSF grants CCF 1422159, 1061938, 0832795 and Simons Collaboration on Algorithms and Geometry grant.

[‡]IBM Research, Bangalore, Karnataka, India. Email: rissaket@in.ibm.com

[1]We conveniently think of the problem as computing the *value* of the optimal labeling.

ISSN 1433-8092

programming (SDP) relaxation. A matching *integrality gap* for this relaxation and its strengthening was shown by Feige and Schechtman [5], and Khot and Vishnoi [9] respectively. Moreover, this approximation factor was shown to be tight by Khot, Kindler, Mossel, and O'Donnell [8][2], assuming Khot's Unique Games Conjecture (UGC) [7]. A similar UGC-tight approximation via an SDP relaxation for the Unique Games problem itself was given by Charikar, Makarychev and Makarychev [2]. Greatly generalizing these results, Raghavendra [16] proved that a certain SDP relaxation achieves an approximation factor arbitrarily close to the optimal for *any* CSP, assuming the UGC. Raghavendra [16] formalized the connection between an integrality gap of the SDP relaxation and the corresponding UGC based hardness factor for a given CSP. For a general k-ary CSP over label set $[L]$, SDP relaxation yields a $O\left(L^k/Lk\right)$-approximation [13], and a corresponding hardness of approximation was recently shown by Chan [1].

While the above line of research underscores the theoretical importance of SDP relaxations, linear programs are usually more efficient in practice and are far more widely used as optimization tools. Thus, it is worthwhile to study how well LP relaxations perform for general classes of problems. In the first such result, Kumar, Manokaran, Tulsiani, and Vishnoi [11] showed a certain LP relaxation to be optimal for a large class of covering and packing problems, assuming the UGC. Dalmau and Krokhin [3] and Kun, O'Donnell, Tamaki, Yoshida, and Zhou [12] independently showed that *width*-1 (see for e.g. [12] for a formal definition) CSPs are *robustly decided* by LP relaxation, i.e. it satisfies almost all hyperedges on an almost satisfiable instance. In recent work, Dalmau, Krokhin, and Manokaran [4] have, assuming the UGC, classified CSPs for which the minimization version[3] admits a constant factor approximation via the LP relaxation.

In this work we study the linear programming analogue of the problem studied by Raghavendra [16], i.e. how well the standard LP relaxation approximates a CSP. We prove the following results.

1.1 Our Results

Let $\mathcal{C}$ be a *class* of constraints and let CSP-$[\mathcal{C}, k, L]$ be the k-ary constraint satisfaction problems over label set $[L]$ where each constraint is from the class $\mathcal{C}$. An instance $\mathcal{I}$ of CSP-$[\mathcal{C}, k, L]$ is a (c, s)-*integrality gap* instance if there is a solution to the LP relaxation $\mathtt{LP}(\mathcal{I})$ given in Figure 1 with objective value at least c, and the optimum of $\mathcal{I}$ is at most s. The main result of this paper is as follows.

Theorem 1.1. *If $\mathcal{I}$ is a (c, s)-integrality gap instance of* CSP-$[\mathcal{C}, k, L]$*, then, assuming the Unique Games Conjecture it is NP-hard to distinguish whether a given instance $\mathcal{H}$ of* CSP-$[\mathcal{C}, k, L]$ *has*

$$\mathsf{opt}(\mathcal{H}) \geq \Omega\left(\frac{c}{k^3 \log L}\right), \quad \text{or} \quad \mathsf{opt}(\mathcal{H}) \leq 4 \cdot s.$$

The LP relaxation in Figure 1 is given by a straightforward relaxation of the integer program for the CSP. The above theorem implies that this basic LP relaxation achieves an approximation factor within a multiplicative $O\left(k^3 \cdot \log L\right)$ of the optimal for any CSP-$[\mathcal{C}, k, L]$, assuming UGC. Note that Raghavendra [16] proved a stronger result: a transformation from a (c, s)-integrality gap for a certain SDP relaxation into a $(c - \varepsilon, s + \varepsilon)$-UGC hardness gap, which implies that the SDP relaxation essentially achieves the optimal approximation. We show that the LP relaxation is nearly as good, i.e. up to a multiplicative loss of $O\left(k^3 \cdot \log L\right)$ in the approximation. Before this work, the best known bound of L^{k-1} was implied by the results of Serna, Trevisan, and Xhafa [17]. In particular, [17] showed an L^{k-1}-approximation for any CSP-$[\mathcal{C}, k, L]$ obtained by the basic LP relaxation, generalizing a previous 2^{k-1}-approximation by Trevisan [18] for the boolean case.

[2][8] also assumed the *Majority is Stablest* conjecture which was later proved by Mossel, O'Donnell, and Oleszkiewicz [15].

[3]The goal in the minimization version of a CSP is to compute a labeling with the minimum number of unsatisfied constraints.

Theorem 1.1 has tight dependence on L: for the Unique Games problem (which is a 2-CSP) on label set $[L]$, the standard LP relaxation has $\Omega(L)$ integrality gap (see Appendix A), whereas a very recent result of Kindler, Kolla, and Trevisan [10] gives an $O(L/\log L)$-approximate SDP rounding algorithm for any 2-CSP over label set $[L]$. The latter improves on a previous $O(L \log\log L/\log L)$-approximate SDP rounding algorithm for Unique Games given in [2].

Our second result pertains to CSPs with a *permutation invariant* set of constraints. Roughly speaking, a set of constraints is permutation invariant if it is closed under the permutation of labels on any of the vertices in the hyperedge. Most of the boolean CSPs such as Max-k-SAT, Max-k-AND, Max-k-XOR etc. are permutation invariant by definition. On larger label sets, Unique Games and Label Cover are well known examples of permutation invariant CSPs. We show that there is a simple randomized LP rounding algorithm such that a weaker version of Theorem 1.1 holds for a corresponding (c, s)*-rounding gap*, which is an instance of a permutation invariant CSP with an LP solution of value c on which the rounding algorithm has an expected payoff at most s. Our rounding algorithm independently rounds each vertex based only on the LP values associated with it. Thus, a *single* constraint suffices to capture its rounding gap. In particular, we prove the following theorem.

Theorem 1.2. *Let $\tilde{\mathcal{I}}$ be a single k-ary hyperedge $\tilde{e}$ with a constraint $C_{\tilde{e}}$ as an instance of a permutation invariant* CSP-$[\mathcal{C}, k, L]$*, which is a (c, s)-rounding gap for the algorithm* Round *given in Figure 2. Then, assuming the Unique Games Conjecture it is NP-hard to distinguish whether a given instance $\mathcal{H}$ of* CSP-$[\mathcal{C}, k, L]$ *has*

$$\mathsf{opt}(\mathcal{H}) \geq \Omega\left(\frac{c}{k^3 \log L}\right), \quad \text{or} \quad \mathsf{opt}(\mathcal{H}) \leq O\left((\log L)^k\right) \cdot s.$$

1.2 Our Techniques

For proving Theorem 1.1, we follow the approach used in earlier works ([16], [11]) of converting an integrality gap instance for the LP relaxation into a UGC-hardness result, which translates the integrality gap into the hardness factor. This reduction essentially involves the construction of a *dictatorship gadget*, which is a toy instance of the CSP-$[\mathcal{C}, k, L]$ distinguishing between "dictator" labelings and "far from dictator" labelings. The construction is illustrated with the following simple example.

Consider an integrality gap instance consisting of just one edge $e = (u, v)$ over label set $[L]$, with the constraint given by the set $C_e \subseteq [L] \times [L]$ of satisfying assignments to (u, v). Let $(\overline{x}, \overline{y})$ be a solution to the corresponding LP relaxation given in Figure 1. It is easy to see that the $\overline{x}$ variables corresponding to u (v) describe a distribution μ_u (μ_v) on $[L]$, and $\overline{y}$ describes a distribution ν_e on $[L] \times [L]$. Furthermore, the marginals of ν_e are μ_u and μ_v. Let $\tilde{\nu}_e = \rho\nu_e + (1 - \rho)(\mu_u \times \mu_v)$, for some parameter ρ. Clearly, the marginals of $\tilde{\nu}_e$ are also μ_u and μ_v.

The vertices of the dictatorship gadget are $\{u, v\} \times [L]^R$ where R is some large enough parameter. The weighted edges are formed as follows. Add an edge between $(u, \overline{r})$ and $(v, \overline{s})$ with weight $\tilde{\nu}_e^R(\overline{r}, \overline{s})$ with the constraint C_e. Here $\tilde{\nu}_e^R$ is the R-wise product distribution of $\tilde{\nu}_e$, i.e. the measure defined by choosing $\overline{r} = (r_1, \ldots, r_R)$ and $\overline{s} = (s_1, \ldots, s_R)$ such that (r_i, s_i) is sampled independently from $\tilde{\nu}_e$, for $i = 1, \ldots, R$.

It is easy to see that for any $i^* = 1, \ldots, R$, over the choice of $\overline{r}$ and $\overline{s}$ above, $(r_{i^*}, s_{i^*}) \in C_e$ with probability at least,

$$\rho \sum_{\overline{\ell} \in C_e} y_{e\overline{\ell}}. \tag{1}$$

Therefore, the above is the fraction of edges in the dictatorship gadget satisfied by labeling each $(u, (r_1, \ldots, r_R))$ with r_{i^*} and each $(v, (s_1, \ldots, s_R))$ with s_{i^*}. More formally, the expression in (1) is the *completeness* of the dictatorship gadget. Note that this is simply ρ times the objective value of the solution $(\overline{x}, \overline{y})$ to $\mathtt{LP}(\mathcal{I})$.

On the other hand, consider a labeling σ to the vertices of the dictatorship gadget. Define functions,

$$f_j(\overline{r}) := \mathbb{1}\{\sigma((u, \overline{r})) = j\}, \qquad g_j(\overline{s}) := \mathbb{1}\{\sigma((v, \overline{s})) = j\}, \tag{2}$$

for $j = 1, \ldots, L$, where $\mathbb{1}\{A\}$ denotes the indicator of the event A. We assume that the labeling σ is "far from dictator", i.e. each of the functions f_j and g_j are far from dictators. Estimating the weighted fraction of edges of the dictatorship gadget satisfied by σ entails analyzing expectations of the form,

$$\mathbb{E}_{\tilde{\nu}_e^R}\left[f_j(\overline{r})g_{j'}(\overline{s})\right], \tag{3}$$

for $1 \leq j, j' \leq L$. In the reduction of Raghavendra [16], such expressions essentially correspond to the payoff yielded by a randomized Gaussian rounding of the SDP solution, under the assumption that σ is far from a dictator. This is obtained by an application of the Invariance Principle developed by Mossel [14]. The parameter ρ is required to be set to only slightly less than 1 in [16] for the application of the Invariance Principle.

In our case the expectation in (3) does not *a priori* correspond to the payoff of any rounding of $(\overline{x}, \overline{y})$. However, we show that setting $\rho \approx (1/\log L)$ is sufficient to ensure,

$$\mathbb{E}_{\tilde{\nu}_e}\left[f_j g_{j'}\right] \approx \mathbb{E}[f_j]\mathbb{E}[g_{j'}], \tag{4}$$

when both $\mathbb{E}[f_j]$ and $\mathbb{E}[g_{j'}]$ are non-negligible. The RHS of the above corresponds to the payoff obtained by assigning u the label j with probability $\mathbb{E}[f_j]$, and independently assigning v label j with probability $\mathbb{E}[g_j]$, $j = 1, \ldots, L$. Thus, the fraction of edges of the dictatorship gadget satisfied by σ, i.e its *soundness*, is essentially bounded by the optimum of the integrality gap instance. There is a $O(\log L)$ loss in the hardness factor, as the completeness decreases due to the setting of ρ.

The proof of Theorem 1.2 proceeds by using a (c, s)-rounding gap $\tilde{\mathcal{I}}$ for the algorithm Round given in Figure 2 to construct a CSP instance, with constraints being permutations of $\tilde{\mathcal{I}}$, which is a $\left(c/4, O\left((\log L)^k\right) \cdot s\right)$-integrality gap for the corresponding LP relaxation. A subsequent application of Theorem 1.1 with this integrality gap instance proves Theorem 1.2.

Organization of the Paper

Theorem 1.1 is restated in Section 3 as Theorem 3.1 which states a hardness reduction from Unique Games. The corresponding Dictatorship Gadget is described in Section 4 and the reduction from Unique Games is given in Section 5. Theorem 3.2, proved in Section 6 gives the transformation from a rounding gap to an integrality gap instance, and along with Theorem 3.1 proves Theorem 1.2.

In the next section we define the constraint satisfaction problem and describe their LP relaxation that we study. The notion of correlated spaces and Gaussian stability bounds used in our reduction and analysis are also described.

2 Preliminaries

We begin by formally defining a constraint satisfaction problem and then describe the LP relaxation that we consider.

2.1 k-ary CSP over label set $[L]$

Let $k \geq 2$ and $L \geq 2$ be positive integers. We say that $C \subseteq [L]^k, C \neq \emptyset$, is a constraint. A collection of such constraints $\mathcal{C}$ is a (k, L)-*constraint class*, i.e.

$$\mathcal{C} \subseteq \left(2^{[L]^k} \setminus \{\emptyset\}\right).$$

We denote by CSP-$[\mathcal{C}, k, L]$ as the class of k-ary constraint satisfaction problems over label set $[L]$, where each constraint is from the class $\mathcal{C}$. Formally, an instance of $\mathcal{I}$ of CSP-$[\mathcal{C}, k, L]$ consists of a finite set of vertices $V_{\mathcal{I}}$, a set of k-uniform ordered hyperedges $E_{\mathcal{I}} \subseteq V_{\mathcal{I}}^k$ and constraints $\{C_e \in \mathcal{C} \mid e \in E\}$. In addition, the hyperedges have normalized weights $\{w_e \geq 0\}_{e \in E_{\mathcal{I}}}$ satisfying $\sum_{e \in E_{\mathcal{I}}} w_e = 1$. A labeling $\sigma : V_{\mathcal{I}} \mapsto [L]$ satisfies the hyperedge $e = (v_1, \ldots, v_k)$ if $(\sigma(v_1), \ldots, \sigma(v_k)) \in C_e$.

As an example, 3-SAT is a constraint satisfaction problem with $k = 3$ over the boolean domain, i.e. $L = 2$. The SAT predicate is over 3 variables. Allowing for negations of the boolean variables yields a constraint class $\mathcal{C}_{3-\mathrm{SAT}}$ consisting of 8 constraints. Each constraint, being an OR over 3 literals, has 7 satisfying assignments (labelings).

Let us denote the weighted fraction of constraints satisfied by any labeling σ by $\mathsf{val}(\mathcal{I}, \sigma)$. The optimum value of the instance is given by,

$$\mathsf{opt}(\mathcal{I}) := \max_{\sigma : V \mapsto [L]} \mathsf{val}(\mathcal{I}, \sigma).$$

2.1.1 Permutation Invariant Constraints

Let $\pi_j : [L] \mapsto [L]$, $j = 1, \ldots, k$, be k permutations. For a constraint $C \subseteq [L]^k$, define the $[\pi_1, \ldots, \pi_k]$-permuted constraint as:

$$[\pi_1, \ldots, \pi_k]C := \{(\pi_1(j_1), \ldots, \pi_k(j_k)) \mid (j_1, \ldots, j_k) \in C\}. \tag{5}$$

A (k, L)-constraint class $\mathcal{C}$ is said to be *permutation invariant* if for every k permutations $\pi_j : [L] \mapsto [L]$ $(1 \leq j \leq k)$, $C \in \mathcal{C}$ implies $[\pi_1, \ldots, \pi_k]C \in \mathcal{C}$. As mentioned earlier, boolean constraint classes such as k-SAT, k-AND and k-XOR are permutation invariant by definition since they are closed under negation of variables. For general L, Unique Games and Label Cover are well studied permutation invariant constraint classes.

2.2 LP Relaxation for CSP-$[\mathcal{C}, k, L]$

The standard linear programming relaxation for an instance $\mathcal{I}$ (as defined above) of CSP-$[\mathcal{C}, k, L]$ is obtained as follows. There is a variable $x_{v\ell}$ for each vertex $v \in V_{\mathcal{I}}$ and label $\ell \in [L]$. For each constraint C_e corresponding to hyperedge $e = (v_1, \ldots, v_k)$, and tuple $\overline{\ell} = (\ell_1, \ldots, \ell_k) \in [L]^k$ of labels, there is a variable $y_{e\overline{\ell}}$. In the integral solution these variables are $\{0, 1\}$-valued denoting the selection the particular label or tuple of labels for the corresponding vertex or hyperedge respectively. To ensure consistency they are appropriately constrained. Allowing the variables to take values in $[0, 1]$, we obtain the LP relaxation denoted by $\mathtt{LP}(\mathcal{I})$ and given in Figure 1.

For a given instance $\mathcal{I}$, let

$$(\overline{x}, \overline{y}) = (\{x_{v\ell}\}_{v \in V_{\mathcal{I}}, \ell \in [L]}, \{y_{e\overline{\ell}}\}_{e \in E_{\mathcal{I}}, \overline{\ell} \in [L]^k}),$$

$$\max \sum_{e \in E_{\mathcal{I}}} w_e \cdot \sum_{\overline{\ell} \in C_e} y_{e\overline{\ell}} \tag{6}$$

subject to,

$$\forall v \in V_{\mathcal{I}}, \qquad \sum_{\ell \in [L]} x_{v\ell} = 1 \tag{7}$$

$$\begin{aligned} &\forall v \in V_{\mathcal{I}} \text{ and,} \\ &e = (v_1, \ldots, v_{i-1}, v, v_{i+1}, \ldots, v_k) \in E_{\mathcal{I}} \text{ and,} \\ &\ell^* \in [L], \qquad \sum_{\overline{\ell} \in [L]^{i-1} \times \{\ell^*\} \times [L]^{k-i}} y_{e\overline{\ell}} = x_{v\ell^*} \end{aligned} \tag{8}$$

$$\forall v \in V_{\mathcal{I}}, \ell \in [L], \qquad x_{v\ell} \geq 0. \tag{9}$$

$$\forall e \in E_{\mathcal{I}}, \overline{\ell} \in [L]^k, \qquad y_{e\overline{\ell}} \geq 0. \tag{10}$$

Figure 1: LP Relaxation $\mathtt{LP}(\mathcal{I})$ for instance $\mathcal{I}$ of CSP-$[\mathcal{C}, k, L]$.

be a valid solution to $\mathtt{LP}(\mathcal{I})$. On this solution, the objective value of the LP is denoted by $\mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y}))$. The *integrality gap*, i.e. how well the LP relaxation approximates the integral optimum on $\mathcal{I}$, is given by,

$$\mathsf{intgap}(\mathcal{I}) := \frac{\mathsf{lpsup}(\mathcal{I})}{\mathsf{opt}(\mathcal{I})}, \tag{11}$$

where,

$$\mathsf{lpsup}(\mathcal{I}) := \sup_{(\overline{x}, \overline{y})} \mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})). \tag{12}$$

A smaller integrality gap – which is always at least 1 – indicates tightness of the LP relaxation. We say that $\mathcal{I}$ is a (c, s)-*integrality gap* instance if,

$$\mathsf{lpsup}(\mathcal{I}) \geq c, \quad \text{and} \quad \mathsf{opt}(\mathcal{I}) \leq s. \tag{13}$$

2.2.1 Smooth LP Solutions

The following shows that the integrality gap is nearly attained by a solution to the LP relaxation which is discrete in the following sense.

Definition 2.1. *Given an instance $\mathcal{I}$ of* CSP-$[\mathcal{C}, k, L]$*, a solution $(\overline{x}, \overline{y})$ to* $\mathtt{LP}(\mathcal{I})$ *is* δ-smooth *if each variable $x_{v\ell}$ is at least δL^{-1} and each variable $y_{e\overline{\ell}}$ is at least δL^{-k}, for any $\delta > 0$.*

The following lemma is proved in Section 10.

Lemma 2.2. *Given an instance $\mathcal{I}$ of* CSP-$[\mathcal{C}, k, L]$*, for any $\delta > 0$ and solution $(\overline{x^*}, \overline{y^*})$ to* $\mathtt{LP}(\mathcal{I})$*, there is an (efficiently computable) δ-smooth solution $(\overline{x}, \overline{y})$ to* $\mathtt{LP}(\mathcal{I})$ *such that,*

$$\mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})) \geq (1 - \delta)\mathsf{lpval}(\mathcal{I}, (\overline{x^*}, \overline{y^*})). \tag{14}$$

In particular, there is a δ-smooth solution $(\overline{x}, \overline{y})$ to $\mathsf{LP}(\mathcal{I})$ such that,

$$\frac{\mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y}))}{\mathsf{opt}(\mathcal{I})} \geq (1-\delta)\mathsf{intgap}(\mathcal{I}). \tag{15}$$

2.3 A Rounding Algorithm for LP

Given an instance $\mathcal{I}$ of CSP-$[\mathcal{C}, k, L]$ and a solution $(\overline{x^*}, \overline{y^*})$ to $\mathsf{LP}(\mathcal{I})$, the rounding algorithm Round is described in Figure 2. The performance of the algorithm is the expected (weighted) fraction of constraints

$\mathsf{Round}(\mathcal{I}, (\overline{x}^*, \overline{y}^*))$:

1. Using Lemma 2.2 compute a 0.1-smooth solution $(\bar{\hat{x}}, \bar{\hat{y}})$ corresponding to $(\overline{x^*}, \overline{y^*})$ satisfying Equation (14).

2. For each vertex $v \in V_{\mathcal{I}}$:

   a. Partition $[L]$ into subsets $\{S_t^v\}_{t=1}^T$, where $S_i^v = \{\ell \in [L] \mid (1/2^t) < \hat{x}_{v\ell} \leq (1/2^{t-1})\}$. Note: $T = O(\log L)$, by 0.1-smoothness of $(\bar{\hat{x}}, \bar{\hat{y}})$.

   b. Choose u.a.r t_v^* from $\{t \mid S_t^v \neq \emptyset\}$.

   c. Label v with ℓ^* chosen u.a.r from $S_{t_v^*}^v$.

Figure 2: Rounding Algorithm for $\mathsf{LP}(\mathcal{I})$ on instance $\mathcal{I}$ of CSP-$[\mathcal{C}, k, L]$.

satisfied by this labeling, and is denoted by $\mathsf{Roundval}(\mathcal{I}, (\overline{x}^*, \overline{y}^*))$. The rounding gap for $\mathcal{I}$ and $(\overline{x}^*, \overline{y^*})$ is given by the following ratio.

$$\mathsf{RoundGap}(\mathcal{I}, (\overline{x^*}, \overline{y^*})) := \frac{\mathsf{lpval}(\mathcal{I}, (\overline{x^*}, \overline{y^*}))}{\mathsf{Roundval}(\mathcal{I}, (\overline{x^*}, \overline{y^*}))}. \tag{16}$$

2.4 Gaussian Stability

We require the following notion of Gaussian stability in our analysis.

Definition 2.3. *Let $\Phi : \mathbb{R} \mapsto [0, 1]$ be the cumulative distribution function of the standard Gaussian. For a parameter ρ, define,*

$$\Gamma_\rho(\mu, \nu) = \Pr[X \leq \Phi^{-1}(\mu), Y \leq \Phi^{-1}(\nu)], \tag{17}$$

where X and Y are two standard Gaussian random variables with covariance matrix $\left(\begin{smallmatrix} 1 & \rho \\ \rho & 1 \end{smallmatrix}\right)$. For $k \geq 3$, $(\rho_1, \ldots, \rho_{k-1}) \in [0, 1]^{k-1}$, and $(\mu_1, \ldots, \mu_k) \in [0, 1]^k$, inductively define,

$$\Gamma_{\rho_1,\ldots,\rho_{k-1}}(\mu_1, \ldots, \mu_k) = \Gamma_{\rho_1}(\mu_1, \Gamma_{\rho_2,\ldots,\rho_{k-1}}(\mu_2, \ldots, \mu_k)). \tag{18}$$

The following key lemma is proved in Section 9.

Lemma 2.4. *Let $k \geq 2$ be an integer and $T \geq 2$ such that $1 \geq \mu_i \geq (1/T)$ for $i = 1, \ldots, k$. Then, there exists a universal constant $C > 0$ such that for any $\varepsilon \in (0, 1/2]$,*

$$\rho = \frac{\varepsilon}{C(k-1)(\log T + \log(1/\varepsilon))}, \tag{19}$$

implies,

$$\Gamma_{\overline{\rho}_{k-1}}(\mu_1, \ldots, \mu_k) \leq (1+\varepsilon)^{k-1} \prod_{i=1}^{k} \mu_i,$$

where $\overline{\rho}_{k-1} = (\rho, \ldots, \rho)$, is a $(k-1)$-tuple with each entry ρ.

2.5 Correlated Spaces

The correlation between two correlated probability spaces is defined as follows.

Definition 2.5. *Suppose $(\Omega^{(1)} \times \Omega^{(2)}, \mu)$ is a finite correlated probability space with the marginal probability spaces $(\Omega^{(1)}, \mu)$ and $(\Omega^{(2)}, \mu)$. The* correlation *between these spaces is,*

$$\rho(\Omega^{(1)}, \Omega^{(2)}; \mu) = \sup\Big\{ |\mathbb{E}_\mu[fg]| \ \mid \ f \in L^2(\Omega^{(1)}, \mu), g \in L^2(\Omega^{(2)}, \mu), \mathbb{E}[f] = \mathbb{E}[g] = 0; \mathbb{E}[f^2], \mathbb{E}[g^2] \leq 1 \Big\}.$$

Let $(\Omega_i^{(1)} \times \Omega_i^{(2)}, \mu_i)_{i=1}^n$ be a sequence of correlated spaces. Then,

$$\rho(\prod_{i=1}^n \Omega_i^{(1)}, \prod_{i=1}^n \Omega_i^{(2)}; \prod_{i=1}^n \mu_i) \leq \max_i \rho(\Omega_i^{(1)}, \Omega_i^{(2)}; \mu_i).$$

Further, the correlation of k correlated spaces $(\prod_{j=1}^k \Omega^{(j)}, \mu)$ is defined as follows:

$$\rho(\Omega^{(1)}, \Omega^{(2)}, \ldots, \Omega^{(k)}; \mu) := \max_{1 \leq i \leq k} \rho\left(\prod_{j=1}^{i-1} \Omega^{(j)} \times \prod_{j=i+1}^{k} \Omega^{(j)}, \Omega^{(i)}; \mu \right).$$

The Bonami-Beckner operator is defined as follows.

Definition 2.6. *Given a probability space (Ω, μ) and $\rho \geq 0$, consider the space $(\Omega \times \Omega, \mu')$ where $\mu'(x, y) = (1-\rho)\mu(x)\mu(y) + \rho \mathbb{1}\{x = y\}\mu(x)$, where $\mathbb{1}\{x = y\} = 1$ if $x = y$ and 0 otherwise. The Bonami-Beckner operator T_ρ is defined by,*

$$(T_\rho f)(x) = \mathbb{E}_{(X,Y) \leftarrow \mu'}\left[f(Y) \mid X = x \right].$$

For product spaces $(\prod_{i=1}^n \Omega_i, \prod_{i=1}^n \mu_i)$, the Bonami-Beckner operator $T_\rho = \otimes_{i=1}^n T_\rho^i$, where T_ρ^i is the operator for the ith space (Ω_i, μ_i).

The influence of a function on a product space is defined as follows.

Definition 2.7. *Let f be a function on $(\prod_{i=1}^n \Omega_i, \prod_{i=1}^n \mu_i)$. The influence of the ith coordinate on f is:*

$$\mathsf{Inf}_i(f) = \mathbb{E}_{\{x_j | j \neq i\}}\left[\mathsf{Var}_{x_i}\left[f(x_1, x_2, \ldots, x_i, \ldots, x_n) \right] \right].$$

The following is a folklore upper bound on the sum of influences of smoothed functions, and is proved as Lemma 1.13 in [19].

Lemma 2.8. *Let f be a function on $(\prod_{i=1}^n \Omega_i, \prod_{i=1}^n \mu_i)$ which takes values in $[-1,1]$. Then,*

$$\sum_{i=1}^{n} \mathsf{Inf}_i(T_{1-\gamma} f) \leq \gamma^{-1}, \tag{20}$$

for any $\gamma \in (0,1]$.

2.6 Useful Invariance and Correlation Bounds

The following key result in Mossel's work [14] shall be used in the analysis of our reduction. We restate Lemma 6.2 of [14].

Lemma 2.9. *Let $(\Omega_1^{(j)}, \ldots, \Omega_n^{(j)})_{j=1}^k$ be k collections of finite probability spaces such that $\{\prod_{j=1}^k \Omega_i^{(j)} \mid i = 1, \ldots, n\}$ are independent. Suppose further that it holds for all $i = 1, \ldots, n$ that $\rho(\Omega_i^{(j)} : 1 \leq j \leq k) \leq \rho$. Then there exists an absolute constant C such that for any $\nu \in (0,1)$,*

$$\gamma = C \frac{(1-\rho)\nu}{\log(1/\nu)},$$

and k functions $\left\{ f_j \in L^2(\prod_{i=1}^n \Omega_i^{(j)}) \right\}_{j=1}^k$, the following holds,

$$\left| \mathbb{E}\left[\prod_{j=1}^k f_j \right] - \mathbb{E}\left[\prod_{j=1}^k T_{1-\gamma} f_j \right] \right| \leq \nu \sum_{j=1}^k \sqrt{\mathrm{Var}[f_j]} \sqrt{\mathrm{Var}\left[\prod_{j'<j} T_{1-\gamma} f_{j'} \prod_{j'>j} f_{j'} \right]}.$$

In particular, if the functions f_j ($1 \leq j \leq k$) take values in $[0,1]$ then,

$$\left| \mathbb{E}\left[\prod_{j=1}^k f_j \right] - \mathbb{E}\left[\prod_{j=1}^k T_{1-\gamma} f_j \right] \right| \leq k\nu.$$

Our analysis shall also utilize the following multi-linear Gaussian stability bound which follows from Theorem 1.14 and Proposition 1.15 of [14] (restated as Theorem 8.1) along with the inductive definition of $\Gamma_{\rho_1, \ldots, \rho_{k-1}}(\mu_1, \ldots, \mu_k)$. A proof is given in Section 8.

Theorem 2.10. *Let $(\prod_{j=1}^k \Omega_i^{(j)}, \mu_i)$ be a sequence of correlated spaces such that for each i, the probability of any atom in $(\prod_{j=1}^k \Omega_i^{(j)}, \mu_i)$ is at least $\alpha \leq 1/2$ and such that $\rho(\Omega_i^{(1)}, \ldots, \Omega_i^{(k)}; \mu_i) \leq \rho$ for all i. Then there exists a universal constant $C > 0$ such that, for every $\nu > 0$, taking*

$$\tau = \left((\nu/k)^{\left(C \frac{k \log(1/\alpha) \log(k/\nu)}{\nu(1-\rho)} \right)} \right) \Big/ k^2 , \tag{21}$$

for functions $\{f_j : \prod_{i=1}^n \Omega_i^{(j)} \mapsto [0,1]\}_{j=1}^k$ that satisfy: ,

$$\forall j, j' \text{ s.t. } 1 \leq j < j' \leq k, \;\; \{i \mid \mathsf{Inf}_i(f_j) > \tau\} \cap \{i \mid \mathsf{Inf}_i(f_{j'}) > \tau\} = \emptyset, \tag{22}$$

the following holds,

$$\mathbb{E}[\prod_{j=1}^k f_j] \leq \Gamma_{\rho, \ldots, \rho}(\mathbb{E}[f_1], \ldots, \mathbb{E}[f_k]) + \nu. \tag{23}$$

2.7 Unique Games Conjecture

UNIQUEGAMES is the following constraint satisfaction problem.

Definition 2.11. *A* UNIQUEGAMES *instance $\mathcal{U}$ consists of a graph $G_{\mathcal{U}} = (V_{\mathcal{U}}, E_{\mathcal{U}})$, a label set $[R]$ and a set of bijections $\{\pi_e : [R] \mapsto [R] \mid e \in E_{\mathcal{U}}\}$. A labeling $\sigma : V_{\mathcal{U}} \mapsto [R]$ satisfies an edge $e = (u, v)$ if $\pi_e(\sigma(v)) = \sigma(u)$. The instance is called d-regular if $G_{\mathcal{U}}$ is d-regular.*

The UNIQUEGAMES problem is: given an instance of UNIQUEGAMES, find an assignment which satisfies the maximum fraction of edges. It is easy to see that if there exists an assignment that satisfies all edges, such an assignment can be efficiently obtained. In other words, the UNIQUEGAMES is easy on satisfiable instances. This is not known to be true for *almost* satisfiable instances, and the following conjecture on the hardness of UNIQUEGAMES on such instances was proposed by Khot [7].

Conjecture 1. *For any constant $\zeta > 0$, there is an integer $R > 0$, such that it is NP-hard, given a regular instance $\mathcal{U}$ of* UNIQUEGAMES *on label set $[R]$, to decide whether,*

YES Case. *There is a labeling to the vertices of $\mathcal{U}$ which satisfies $(1 - \zeta)$ fraction of its edges.*

NO Case. *Any labeling satisfies at most ζ fraction of the edges.*

3 Our Results restated

The following is a restatement of Theorem 1.1 as a hardness reduction from UNIQUEGAMES.

Theorem 3.1. *Let $k \geq 2$ and $L \geq 2$ be positive integers. Let $\mathcal{I}$ be a (c, s)-integrality gap instance of* CSP-$[\mathcal{C}, k, L]$. *Then, there is a reduction from an instance $\mathcal{U}$ of* UNIQUEGAMES *given by Conjecture 1 with a small enough parameter ζ, to an instance $\mathcal{H}$ of* CSP-$[\mathcal{C}, k, L]$ *such that,*

YES Case. *If $\mathcal{U}$ is a YES instance, then*

$$\mathsf{opt}(\mathcal{H}) \geq \Omega\left(\frac{c}{k^3 \log L}\right).$$

NO Case. *If $\mathcal{U}$ is a NO instance, then,*

$$\mathsf{opt}(\mathcal{H}) \leq 4 \cdot s.$$

Theorem 3.1 is obtained by combining the *dictatorship gadget* constructed in Section 4 with the hard instance of UNIQUEGAMES. As the name suggests, this gadget distinguishes between labelings defined by a dictator and those which are not. The dictatorship gadget illustrates the main ideas of the hardness reduction and is derived from the integrality gap instance $\mathcal{I}$ of CSP-$[\mathcal{C}, k, L]$, and is also a CSP-$[\mathcal{C}, k, L]$ instance. This notion is the same as defined by Raghavendra [16] and can be converted into a hardness reduction from UNIQUEGAMES using techniques from Section 6 of [16]. However, to avoid describing the framework of [16] in detail, we provide a direct hardness reduction proving Theorem 3.1 in Section 5.

Our second result Theorem 1.2 is implied by the following theorem and an application of Theorem 3.1.

Theorem 3.2. *Let $k \geq 2$ and $L \geq 2$ be positive integers. Let $\tilde{\mathcal{I}}$ be an instance of* CSP-$[\mathcal{C}, k, L]$ *consisting of one hyperedge $\tilde{e}$ and its constraint $C_{\tilde{e}}$, and $(\overline{x^*}, \overline{y^*})$ be a solution to $\mathsf{LP}(\tilde{\mathcal{I}})$ such that,*

$$\mathsf{lpval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*})) \geq \mathsf{Roundval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*})). \tag{24}$$

Then, there exists an instance $\mathcal{I}$ whose size depends only on L and k with constraints which are permutations of $C_{\tilde{e}}$, and a solution $(\overline{x}, \overline{y})$ to $\mathtt{LP}(\mathcal{I})$ *such that,*

$$\mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})) \geq \frac{\mathsf{lpval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))}{4}, \tag{25}$$

and,

$$\mathsf{opt}(\mathcal{I}) \leq O\left((\log L)^k\right) \mathsf{Roundval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*})). \tag{26}$$

Theorem 3.2 is proved in Section 6.

4 Dictatorship Gadget

We begin with the description of some probability spaces defined using solutions to the LP relaxation given in Figure 1.

4.1 Probability Spaces given by solutions to LP

For a CSP-$[\mathcal{C}, k, L]$ instance $\mathcal{I}$ and a valid solution $(\overline{x}, \overline{y})$ to $\mathtt{LP}(\mathcal{I})$, we define the following useful probability spaces. For each $v \in V_{\mathcal{I}}$, let μ_v be a probability measure over $[L]$ defined as:

$$\mu_v(\ell) = x_{v\ell}, \;\; \forall \ell \in [L]. \tag{27}$$

Also, define for each hyperedge $e = (v_1, \ldots, v_k) \in E_{\mathcal{I}}$, a probability measure ν_e over $[L]^k$ as:

$$\nu_e(\overline{\ell}) = y_{e\overline{\ell}}, \;\; \forall \overline{\ell} \in [L]^k. \tag{28}$$

For a parameter $\rho \in [0, 1]$ define,

$$\tilde{\nu}_{e\rho} = \rho \nu_e + (1 - \rho) \prod_{i=1}^{k} \mu_{v_i}. \tag{29}$$

Therefore, $\nu_e = \tilde{\nu}_{e\rho}$ for $\rho = 1$. Since $(\overline{x}, \overline{y})$ is a valid solution, it is easy to see that for a hyperedge e and its ith vertex v, the marginal distribution of ν_e at the ith coordinate is same as the distribution μ_v. The same is true for $\tilde{\nu}_{e\rho}$ for any $\rho \in [0, 1]$. Also, in the notation of Mossel [14], for the probability space $(\prod_{i=1}^{k}[L]; \tilde{\nu}_{e\rho})$,

$$\rho([L], \ldots, [L]; \tilde{\nu}_{e\rho}) \leq \rho, \tag{30}$$

where $\rho([L], \ldots, [L]; \tilde{\nu}_{e\rho})$ is the correlation of the probability space $(\prod_{i=1}^{k}[L]; \tilde{\nu}_{e\rho})$. The above follows from the definition of $\tilde{\nu}_{e\rho}$.

Further, we denote by $\tilde{\nu}_{e\rho}^{R}$ the product measure on $([L]^R)^k$, defined as:

$$\tilde{\nu}_{e\rho}^{R}(\overline{r}^1, \ldots, \overline{r}^k) = \prod_{i=1}^{R} \tilde{\nu}_{e\rho}(r_i^1, \ldots, r_i^k), \tag{31}$$

where $\overline{r}^j = (r_1^j, \ldots, r_R^j) \in [L]^R$ for $j = 1, \ldots, k$.

4.2 Gadget Construction

Let $\mathcal{I}$ be a CSP-$[\mathcal{C}, k, L]$ instance. From Lemma 2.2, let $(\overline{x}, \overline{y})$ be a δ-smooth solution to $\mathtt{LP}(\mathcal{I})$ satisfying Equation (15) for a parameter $\delta \in [0, 1]$.

The dictatorship gadget is parametrized by a large enough positive integer R and a correlation $\rho \in [0, 1]$ to be set later. We denote the gadget by $\mathcal{D}$ and its set of vertices and hyperedges as $V_{\mathcal{D}}$ and $H_{\mathcal{D}}$ respectively. Each hyperedge $\hat{e} \in E_{\mathcal{D}}$ has a constraint $C_{\hat{e}}$ from the class $\mathcal{C}$ and a normalized positive weight $w_{\hat{e}}$.

Vertices. $V_{\mathcal{D}} := V_{\mathcal{I}} \times [L]^R$. Denote by $V_{\mathcal{D}}^v$ the set of vertices $\{(v, \overline{r}) \mid \overline{r} \in [L]^R\}$ for $v \in V_{\mathcal{I}}$. Thus, $V_{\mathcal{D}} = \cup_{v \in V_{\mathcal{I}}} V_{\mathcal{D}}^v$.

Hyperedges. Let $e = (v_1, \ldots, v_k) \in E_{\mathcal{I}}$. For any $(\overline{r}_1, \ldots, \overline{r}_k) \in ([L]^R)^k$ there is a hyperedge $\hat{e} = ((v_1, \overline{r}_1), \ldots, (v_k, \overline{r}_k))$ in $E_{\mathcal{D}}$, with $C_{\hat{e}} = C_e$. The weight $w_{\hat{e}}$ is given by,

$$w_{\hat{e}} = w_e \cdot \tilde{\nu}_{e\rho}^R(\overline{r}_1, \ldots, \overline{r}_k). \tag{32}$$

It is easy to see that $w_{\hat{e}}$ is a normalized weight function. For convenience, let $E_{\mathcal{D}}(e)$ be the set of hyperedges in $\mathcal{D}$ corresponding to $e \in E_{\mathcal{I}}$.

The above completes the description of the dictatorship gadget $\mathcal{D}$. The gadget distinguishes between dictator labelings and labelings far from a dictator, as shown in the YES and NO cases below.

4.3 YES Case

Let us fix $i^* \in [R]$. Define a labeling σ^* to $V_{\mathcal{D}}$ where,

$$\sigma^*((v, (r_1, \ldots, r_R))) = r_{i^*}, \tag{33}$$

for each $v \in V_{\mathcal{I}}$ and $(r_1, \ldots, r_R) \in [L]^R$. The following lemma shows that σ^* is a good labeling.

Lemma 4.1. *For σ^* defined as above,*

$$\mathsf{val}(\mathcal{D}, \sigma^*) \geq \rho \cdot \mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})).$$

Proof. Consider any hyperedge $e = (v_1, \ldots, v_k) \in E_{\mathcal{I}}$. The (weighted) fraction of hyperedges in $E_{\mathcal{D}}(e)$ satisfied by σ^* is given by,

$$\begin{aligned} &\sum_{(\overline{r}_1, \ldots, \overline{r}_k) \in ([L]^R)^k} w_e \cdot \tilde{\nu}_{e\rho}^R(\overline{r}_1, \ldots, \overline{r}_k) \mathbb{1}\{(\sigma^*((v_1, \overline{r}_1)), \ldots \sigma^*((v_k, \overline{r}_k))) \in C_e\}, \\ = &\sum_{(\overline{r}_1, \ldots, \overline{r}_k) \in ([L]^R)^k} w_e \cdot \tilde{\nu}_{e\rho}^R(\overline{r}_1, \ldots, \overline{r}_k) \mathbb{1}\{(\overline{r}_1(i^*), \ldots \overline{r}_k(i^*)) \in C_e\}, \end{aligned} \tag{34}$$

where $\overline{r}_j(i)$ is the ith coordinate of $\overline{r}_j$, $\forall j = 1, \ldots, k$. Since $(\overline{r}_1(i), \ldots \overline{r}_k(i))$ is independently chosen for $i = 1, \ldots, R$, the RHS of Equation (34) can be rewritten as,

$$w_e \cdot \mathbb{E}_{(r_1, \ldots, r_k) \in_{\tilde{\nu}_{e\rho}} [L]^k} \left[\mathbb{1}\{(r_1, \ldots, r_k) \in C_e\}\right] \geq \rho \cdot w_e \cdot \sum_{\overline{\ell} \in C_e} y_{e\overline{\ell}}, \tag{35}$$

where inequality follows from the definition of $\tilde{\nu}_{e\rho}$. Therefore,

$$\mathsf{val}(\mathcal{D}, \sigma^*) \geq \sum_{e \in E_{\mathcal{I}}} w_e \cdot \rho \sum_{\overline{\ell} \in C_e} y_{e\overline{\ell}} = \rho \cdot \mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})).$$

□

4.4 NO Case

Let σ be a labeling to $V_{\mathcal{D}}$. For any $v \in V_I$ define functions $f_\ell^v : [L]^R \mapsto [0,1]$ for all $\ell \in [L]$ as,

$$f_\ell^v(\overline{r}) := \mathbb{1}\{\sigma(v,\overline{r}) = \ell\}. \tag{36}$$

It follows that,

$$\mathbb{E}[f_\ell^v] \in [0,1], \tag{37}$$

and,

$$\sum_{\ell\in[L]} \mathbb{E}[f_\ell^v] = 1, \tag{38}$$

where the expectation is over the product measure μ_v^R. We now set the parameter ρ in the construction of the dictatorship gadget as follows:

$$\rho := \frac{1}{C(k-1)k[k\log L + \log(2/\varepsilon) + \log k]}, \tag{39}$$

where C is the constant from Lemma 2.4 and $\varepsilon \in [0,1]$ is a parameter. The following lemma gives an upper bound on the value achieved by a non-dictator labeling σ.

Lemma 4.2. *For every $\varepsilon > 0$, there is a constant $\tau > 0$ depending only on ε, L, k and δ such that the following holds. Suppose that for any two vertices $u, v \in V_{\mathcal{I}}$ and labels $\ell, \ell' \in [L]$,*

$$\{i \in [R] \mid \mathsf{Inf}_i(f_\ell^u) > \tau\} \cap \{i \in [R] \mid \mathsf{Inf}_i(f_{\ell'}^v) > \tau\} = \emptyset. \tag{40}$$

Then,

$$\mathsf{val}(\mathcal{D},\sigma) \le 3\cdot \mathsf{opt}(\mathcal{I}) + \varepsilon. \tag{41}$$

Proof. For any hyperedge $e = (v_1,\ldots,v_k) \in E_{\mathcal{I}}$, the fraction of edges in $E_{\mathcal{D}}(e)$ satisfied by σ is,

$$\begin{aligned} &\mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)\leftarrow\tilde{\nu}_{e\rho}^R}\left[\mathbb{1}\{(\sigma((v_1,\overline{r}_1)),\ldots\sigma((v_k,\overline{r}_k))) \in C_e\}\right], \\ = \ &\mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}\left[\sum_{(\ell_1,\ldots,\ell_k)\in C_e}\prod_{j=1}^k f_{\ell_j}^{v_j}(\overline{r}_j)\right] = \sum_{(\ell_1,\ldots,\ell_k)\in C_e}\mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}\left[\prod_{j=1}^k f_{\ell_j}^{v_j}(\overline{r}_j)\right]. \end{aligned} \tag{42}$$

Consider any $f_{\ell_j}^{v_j}$ such that $\mathbb{E}[f_{\ell_j}^{v_j}] \le (\varepsilon/2)L^{-k}$. Call any expectation of products on the RHS of Equation (42) in which $f_{\ell_j}^{v_j}$ occurs a *light* expectation. Any light expectation is also bounded by $(\varepsilon/2)L^{-k}$. Since, there are at most L^k expectations in the sum, one can ignore all light expectations on the RHS, losing only an additive factor of $(\varepsilon/2)$ in the upper bound. The remaining expectations are called *heavy* and are analyzed as follows.

Since $(\overline{x},\overline{y})$ is a δ-smooth solution, the construction of the probability space $([L]^k;\tilde{\nu}_{e\rho})$ implies that measure of its smallest atom is at least $(1-\rho)\left(\delta L^{-1}\right)^k$. The correlation of this space is also at most ρ. Our setting of ρ depends only on ε, L and k. Thus, assuming the supposition in the statement of the lemma for a τ that depends only on L, k, ε and δ, one can apply Theorem 2.10 to obtain,

$$\mathbb{E}\left[\prod_{j=1}^k f_{\ell_j}^{v_j}\right] \le \Gamma_{\overline{\rho}_{k-1}}\left(\mathbb{E}\left[f_{\ell_1}^{v_1}\right],\ldots,\mathbb{E}\left[f_{\ell_k}^{v_k}\right]\right) + (\varepsilon/2)L^{-k}, \tag{43}$$

where $\overline{\rho}_{k-1} = (\rho, \ldots, \rho)$ is a $(k-1)$-tuple with each entry ρ. Since we assume that all the expectations in the RHS of the above are at least $(\varepsilon/2)L^{-k}$, by our setting of ρ and Lemma 2.4,

$$\Gamma_{\overline{\rho}_{k-1}}\left(\mathbb{E}\left[f^{v_1}_{\ell_1}\right], \ldots, \mathbb{E}\left[f^{v_k}_{\ell_k}\right]\right) \leq \left(1 + \frac{1}{k}\right)^{k-1} \prod_{j=1}^{k} \mathbb{E}\left[f^{v_j}_{\ell_j}\right] \leq 3 \cdot \prod_{j=1}^{k} \mathbb{E}\left[f^{v_j}_{\ell_j}\right]. \tag{44}$$

Combining the above with Equation (43), we obtain that for the heavy expectations on the RHS of Equation (42),

$$\mathbb{E}\left[\prod_{j=1}^{k} f^{v_j}_{\ell_j}\right] \leq 3 \cdot \prod_{j=1}^{k} \mathbb{E}\left[f^{v_j}_{\ell_j}\right] + (\varepsilon/2)L^{-k}. \tag{45}$$

Substituting the above into Equation (42), along with the above observation that the sum of the light expectations is at most $(\varepsilon/2)$, we obtain that the fraction of edges in $E_{\mathcal{D}}(e)$ satisfied by σ is at most,

$$3 \cdot \sum_{(\ell_1,\ldots,\ell_k)\in C_e} \prod_{j=1}^{k} \mathbb{E}\left[f^{v_j}_{\ell_j}\right] + \varepsilon. \tag{46}$$

The sum in the above expression is simply the probability that the hyperedge $e \in \mathbb{E}_{\mathcal{I}}$ is satisfied when every vertex v is independently assigned a label ℓ with probability $\mathbb{E}[f^{v}_{\ell}]$. Taking a weighted sum over all $e \in E_{\mathcal{I}}$ yields the expected value of this assignment which is at most $\mathsf{opt}(\mathcal{I})$. This completes the proof. $\square$

5 Hardness Reduction from UNIQUEGAMES

The hardness reduction essentially combines a hard instance of UNIQUEGAMES with the dictatorship gadget constructed in Section 4. We first give the reduction which parametrized by $\varepsilon, \delta, \rho \in [0, 1]$ to be set later. This is followed by the analysis of the YES and NO cases, and finally we show that an appropriate setting of the parameters in the reduction implies Theorem 3.1.

As in Section 4, $\mathcal{I}$ is a CSP-$[\mathcal{C}, k, L]$ instance and let $(\overline{x}, \overline{y})$ be a δ-smooth solution to $\mathsf{LP}(\mathcal{I})$ satisfying Equation (15). Let $\mathcal{U}(G_{\mathcal{U}} = (V_{\mathcal{U}}, E_{\mathcal{U}}), [R], \{\pi_e\}_{e\in E_{\mathcal{U}}})$ be a d-regular instance of UNIQUEGAMES with parameter $\zeta > 0$ (to be chosen later) as given in Conjecture 1.

The hardness reduction produces an instance $\mathcal{H}$ of CSP-$[\mathcal{C}, k, L]$ with $V_{\mathcal{H}}$ and $E_{\mathcal{H}}$ as its vertices and hyperedges respectively. Each hyperedge $\tilde{e} \in E_{\mathcal{H}}$ has a constraint $C_{\tilde{e}}$ from the class $\mathcal{C}$ and a normalized positive weight $w_{\tilde{e}}$.

Vertices. $V_{\mathcal{H}} := V_{\mathcal{U}} \times V_{\mathcal{I}} \times [L]^R$. Denote by $V_{\mathcal{H}}(\hat{u}, v)$ the set of vertices $\{(\hat{u}, v, \overline{r}) \mid \overline{r} \in [L]^R\}$ for $\hat{u} \in V_{\mathcal{U}}$ and $v \in V_{\mathcal{I}}$. Thus, $V_{\mathcal{H}} = \cup_{\hat{u}\in V_{\mathcal{U}}} \cup_{v\in V_{\mathcal{I}}} V_{\mathcal{H}}(\hat{u}, v)$.

Hyperedges. For convenience we define the following notation. For a bijection $\pi : [R] \mapsto [R]$ and $\overline{r} \in [L]^R$, let $(\overline{r} \circ \pi) \in [L]^R$ where,

$$(\overline{r} \circ \pi)(i) = \overline{r}(\pi(i)), \quad \forall i \in [R]. \tag{47}$$

The hyperedges are constructed as follows. Let $\hat{u} \in V_{\mathcal{U}}$ and let $(\hat{v}_1, \ldots, \hat{v}_k)$ be a k-tuple of its neighbors in $G_{\mathcal{U}}$ via edges $\hat{e}_j = (\hat{u}, \hat{v}_j)$, $j = 1, \ldots, k$. For each $\hat{u}$ there are d^k such tuples. Let $e = (v_1, \ldots, v_k) \in E_{\mathcal{I}}$. For any $(\overline{r}_1, \ldots, \overline{r}_k) \in ([L]^R)^k$ there is a hyperedge $\tilde{e} = ((\hat{v}_1, v_1, (\overline{r}_1 \circ \pi_{\hat{e}_1})), \ldots, (\hat{v}_k, v_k, (\overline{r}_k \circ \pi_{\hat{e}_k})))$ in $E_{\mathcal{H}}$, with $C_{\tilde{e}} = C_e$. The weight $w_{\tilde{e}}$ is given by,

$$w_{\tilde{e}} = \left(\frac{1}{d^k\,|V_{\mathcal{U}}|}\right) \cdot w_e \cdot \tilde{\nu}^{R}_{e\rho}(\overline{r}_1, \ldots, \overline{r}_k). \tag{48}$$

Observe that there are $d^k|V_{\mathcal{U}}|$ choices of $\hat{u}$ and a k-tuple of its neighbors. Therefore, $w_{\tilde{e}}$ is a product of three independent probability measures, and is thus a normalized weight function. For convenience, let $E_{\mathcal{H}}(\hat{u}, (\hat{v}_1, \ldots, \hat{v}_k), e)$ be the set of hyperedges in $\mathcal{H}$ corresponding to $\hat{u} \in V_{\mathcal{U}}$, the k-tuple $(\hat{v}_1, \ldots, \hat{v}_k)$ of its neighbors, and $e \in E_{\mathcal{I}}$.

The above completes the construction of the instance $\mathcal{H}$.

5.1 YES Case

Let $\hat{\sigma}$ be a labeling to the vertices of $\mathcal{U}$ from the set $[R]$ that satisfies $(1-\zeta)$ fraction of edges. Define a labeling σ^* to $V_{\mathcal{H}}$ where,

$$\sigma^*((\hat{u}, v, (r_1, \ldots, r_R))) = r_{\hat{\sigma}(\hat{u})}, \tag{49}$$

for each $\hat{u} \in V_{\mathcal{U}}$, $v \in V_{\mathcal{I}}$, and $(r_1, \ldots, r_R) \in [L]^R$. The following lemma shows that σ^* is a good labeling.

Lemma 5.1. *For σ^* defined as above,*

$$\mathsf{val}(\mathcal{H}, \sigma^*) \geq \rho \cdot \mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})) - k\zeta.$$

Proof. Since $\hat{\sigma}$ satisfies at least $(1-\zeta)$ fraction of edges, the fraction of choices of $\hat{u}$ and a k-tuple of its neighbors $(\hat{v}_1, \ldots, \hat{v}_k)$ such all of the edges $\hat{e}_j = (\hat{u}, \hat{v}_j)$ $(1 \leq j \leq k)$ are satisfied by $\hat{\sigma}$ is at least $(1 - k\zeta)$. Thus, losing an additive factor of $k\zeta$ we assume this to be true for a fixed choice of $\hat{u}$ and a k-tuple of its neighbors $(\hat{v}_1, \ldots, \hat{v}_k)$.

Consider any hyperedge $e = (v_1, \ldots, v_k) \in E_{\mathcal{I}}$. The (weighted) fraction of hyperedges in $E_{\mathcal{H}}(\hat{u}, (\hat{v}_1, \ldots, \hat{v}_k), e)$ satisfied by σ^* is given by,

$$\begin{aligned} & \sum_{(\overline{r}_1, \ldots, \overline{r}_k) \in ([L]^R)^k} w_e \cdot \tilde{\nu}_{e\rho}^R(\overline{r}_1, \ldots, \overline{r}_k) \mathbb{1}\big\{(\sigma^*((\hat{v}_1, v_1, (\overline{r}_1 \circ \pi_{\hat{e}_1}))), \ldots, \sigma^*((\hat{v}_k, v_k, (\overline{r}_k \circ \pi_{\hat{e}_k})))) \in C_e\big\} \\ = & \sum_{(\overline{r}_1, \ldots, \overline{r}_k) \in ([L]^R)^k} w_e \cdot \tilde{\nu}_{e\rho}^R(\overline{r}_1, \ldots, \overline{r}_k) \mathbb{1}\big\{((\overline{r}_1 \circ \pi_{\hat{e}_1})(\hat{\sigma}(\hat{v}_1)), \ldots, (\overline{r}_k \circ \pi_{\hat{e}_k})(\hat{\sigma}(\hat{v}_k))) \in C_e\big\}, \end{aligned} \tag{50}$$

where $(\overline{r}_j \circ \pi_{\hat{e}_j})(i)$ is the ith coordinate of $(\overline{r}_j \circ \pi_{\hat{e}_j})$, $\forall j = 1, \ldots, k$. Observe that,

$$(\overline{r}_j \circ \pi_{\hat{e}_j})(\hat{\sigma}(\hat{v}_j)) = \overline{r}_j \left(\pi_{\hat{e}_j}(\hat{\sigma}(\hat{v}_j))\right) = \overline{r}_j(\hat{\sigma}(\hat{u})),$$

since $\hat{\sigma}$ satisfies all the edges $\hat{e}_j = (\hat{u}, \hat{v}_j)$ $(1 \leq j \leq k)$. Also, $(\overline{r}_1(i), \ldots \overline{r}_k(i))$ is independently chosen for $i = 1, \ldots, R$. Therefore, the RHS of Equation (50) can be rewritten as,

$$w_e \cdot \mathbb{E}_{(r_1, \ldots, r_k) \in_{\tilde{\nu}_{e\rho}} [L]^k} \left[\mathbb{1}\{(r_1, \ldots, r_k) \in C_e\}\right] \geq \rho \cdot w_e \cdot \sum_{\overline{\ell} \in C_e} y_{e\overline{\ell}}, \tag{51}$$

where inequality follows from the definition of $\tilde{\nu}_{e\rho}$. Summed over all edges $e \in E_{\mathcal{I}}$, we obtain that the fraction of edges corresponding to the our choice of $\hat{u}$ and $(\hat{v}_1, \ldots, \hat{v}_k)$ is at least,

$$\sum_{e \in E_{\mathcal{I}}} \left[w_e \cdot \rho \sum_{\overline{\ell} \in C_e} y_{e\overline{\ell}}\right] = \rho \cdot \mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})).$$

Combining the above with the additive loss of $k\zeta$ incurred towards the choice of $\hat{u}$ and $(\hat{v}_1, \ldots, \hat{v}_k)$, we obtain,

$$\mathsf{val}(\mathcal{H}, \sigma^*) \geq \rho \cdot \mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})) - k\zeta.$$

$\square$

5.2 NO Case

Let σ be any labeling to $V_{\mathcal{D}}$. For any $\hat{v} \in V_{\mathcal{U}}$ and $v \in V_I$ define functions $f_\ell^{\hat{v}v} : [L]^R \mapsto [0,1]$ for all $\ell \in [L]$ as,

$$f_\ell^{\hat{v}v}(\overline{r}) := \mathbb{1}\{\sigma(\hat{v}, v, \overline{r}) = \ell\}, \tag{52}$$

so that,

$$\mathbb{E}[f_\ell^{\hat{v}v}] \in [0,1], \tag{53}$$

and,

$$\sum_{\ell \in [L]} \mathbb{E}[f_\ell^{\hat{v}v}] = 1, \tag{54}$$

where the expectation is over the product measure μ_v^R. For $\gamma \in [0,1]$, let $T_{1-\gamma}$ be the Bonami-Beckner operator from Definition 2.6. Given any $\hat{v} \in V_{\mathcal{U}}$ define,

$$S_{\hat{v}} := \bigcup_{\substack{v \in V_{\mathcal{I}} \\ \ell \in [L]}} \{i \in [R] \mid \mathsf{Inf}_i(T_{1-\gamma} f_\ell^{\hat{v}v}) > \tau\} \tag{55}$$

By Lemma 2.8,

$$\sum_{i=1}^{R} \mathsf{Inf}_i(T_{1-\gamma} f_\ell^{\hat{v}v}) \leq 1/\gamma, \tag{56}$$

for any $\hat{v} \in V_{\mathcal{U}}$, $v \in V_{\mathcal{I}}$ and $\ell \in [L]$. Therefore,

$$|S_{\hat{v}}| < \frac{|V_{\mathcal{I}}| \cdot L}{\tau\gamma}. \tag{57}$$

The following lemma essentially bounds the probability of pairs of vertices with common influential coordinates.

Lemma 5.2. *Let $\hat{u}$ be a vertex chosen u.a.r from $V_{\mathcal{U}}$, and $(\hat{v}_1, \ldots, \hat{v}_k)$ a k-tuple of its neighbors chosen u.a.r. Let,*

$$\eta := \Pr_{\hat{u},(\hat{v}_1,\ldots,\hat{v}_k)} \left[\exists j, j' \text{ s.t. } 1 \leq j < j' \leq k \text{ and, } \pi_{\hat{e}_j}(S_{\hat{v}_j}) \cap \pi_{\hat{e}_{j'}}(S_{\hat{v}_{j'}}) \neq \emptyset\right],$$

where $\pi(S) = \{\pi(s) | s \in S\}$. Then,

$$\eta \leq \frac{4\zeta\left(k^2 L^2 |V_{\mathcal{I}}|^2\right)}{(\gamma\tau)^2}. \tag{58}$$

Proof. Consider the following randomized labeling for the vertices $\hat{u} \in V_{\mathcal{I}}$: with probability $1/2$ choose a uniformly random label from $S_{\hat{u}}$; and with probability $1/2$ choose a uniformly random neighbor $\hat{v}$ of $\hat{u}$ and choose a label uniformly at random from $\pi_{\hat{e}}(S_{\hat{v}})$, where $\hat{e} = (\hat{u}, \hat{v})$. Now, consider the probability that this labeling satisfies an edge $(\hat{u}, \hat{v}')$ chosen uniformly at random. By regularity of $G_{\mathcal{U}}$, this is same as choosing a vertex $\hat{u}$ uniformly at random, and choosing one of its neighbors $\hat{v}'$ uniformly at random. The neighbor $\hat{v}$ used to define the randomized labeling is another of its neighbors chosen independently and uniformly at random. Thus, $\hat{v}$ and $\hat{v}'$ can be thought of as the jth and j'th coordinates of a uniformly random k-tuple of neighbors of $\hat{u}$, for a uniformly random choice j and j' of indices such that $1 \leq j < j' \leq 1$. From the assumption of the lemma, over the choice of $\hat{u}$, $\hat{v}$ and $\hat{v}'$, with probability at least $\eta/(k^2)$,

$$\pi_{\hat{e}}(S_{\hat{v}}) \cap \pi_{\hat{e}'}(S_{\hat{v}'}) \neq \emptyset \tag{59}$$

With a further probability of $1/4$, $\hat{u}$ is labeled by a uniformly chosen label from $\pi_{\hat{e}}(S_{\hat{v}})$, and $\hat{v}'$ is labeled by a uniformly chosen label from $S_{\hat{v}'}$. By the condition in Equation (59), with a further probability of,

$$\frac{1}{|S_{\hat{v}}|\,|S_{\hat{v}'}|},$$

this choice satisfies satisfies $(\hat{u}, \hat{v}')$. Using Equation (57), this implies that the expected fraction of edges satisfied is at least,

$$\frac{\eta(\gamma\tau)^2}{4\left(k^2 L^2 |V_{\mathcal{I}}|^2\right)}.$$

The above is at most ζ and substituting for it proves the lemma. $\square$

The parameter ρ is set as follows:

$$\rho := \frac{1}{C(k-1)k[k\log L + \log(4/\varepsilon) + \log k]}, \tag{60}$$

where C is the constant from Lemma 2.4. The following is the main lemma showing the upper bound on the optimum in the NO case.

Lemma 5.3. *For the above setting of ρ, and sufficiently small choice of $\zeta > 0$,*

$$\mathsf{val}(\mathcal{H}, \sigma) \leq 3 \cdot \mathsf{opt}(\mathcal{I}) + \varepsilon.$$

Proof. For a choice of parameters $\gamma, \tau > 0$ (which we shall set later) let η be as given in Lemma 5.2. By averaging we may assume that for at least $(1 - \sqrt{\eta})$ fraction of the vertices $\hat{u} \in V_{\mathcal{U}}$, for $(1 - \sqrt{\eta})$ fraction of choices of the k-tuple of its neighbors $(\hat{v}_1, \ldots, \hat{v}_k)$,

$$\forall 1 \leq j < j' \leq k, \pi_{\hat{e}_j}(S_{\hat{v}_j}) \cap \pi_{\hat{e}_{j'}}(S_{\hat{v}_{j'}}) = \emptyset. \tag{61}$$

We refer to such vertices $\hat{u}$ as *good*, and the k-tuples of its neighbors $(\hat{v}_1, \ldots, \hat{v}_k)$ satisfying Equation (61) as its *good* k-tuples. Note that the condition in Equation (61) depends on γ and τ. We have the following intermediate lemma.

Lemma 5.4. *For a sufficiently small choice of γ and τ depending on L, k, δ and ε the following holds. For every choice of a good vertex $\hat{u}$ and a good k-tuple $(\hat{v}_1, \ldots, \hat{v}_k)$ of its neighbors, the fraction of hyperedges in $E_{\mathcal{H}}$ corresponding to the choice of $\hat{u}$ and $(\hat{v}_1, \ldots, \hat{v}_k)$ satisfied by σ is at most,*

$$3 \cdot \sum_{e=(v_1,\ldots,v_k)\in E_{\mathcal{I}}} w_e \cdot \left[\sum_{(\ell_1,\ldots,\ell_k)\in C_e} \prod_{j=1}^{k} \mathbb{E}\left[f_{\ell_j}^{\hat{v}_j v_j}\right]\right] + (3\varepsilon/4). \tag{62}$$

Proof. Fix a hyperedge $e = (v_1, \ldots, v_k) \in E_{\mathcal{I}}$. The fraction of hyperedges in $E_{\mathcal{H}}(\hat{u}, (\hat{v}_1, \ldots, \hat{v}_k), e)$ satisfied by σ is,

$$\begin{aligned}
&\mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)\leftarrow\tilde{\nu}_{e\rho}^R}\left[\mathbb{1}\{(\sigma((\hat{v}_1, v_1, (\overline{r}_1 \circ \pi_{\hat{e}_1}))), \ldots \sigma((\hat{v}_k, v_k, (\overline{r}_k \circ \pi_{\hat{e}_k})))) \in C_e\}\right], \\
= \;& \mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}\left[\sum_{(\ell_1,\ldots,\ell_k)\in C_e} \prod_{j=1}^{k} f_{\ell_j}^{\hat{v}_j v_j}((\overline{r}_j \circ \pi_{\hat{e}_j}))\right], \\
= \;& \sum_{(\ell_1,\ldots,\ell_k)\in C_e} \mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}\left[\prod_{j=1}^{k} f_{\ell_j}^{\hat{v}_j v_j}((\overline{r}_j \circ \pi_{\hat{e}_j}))\right].
\end{aligned} \tag{63}$$

Consider any $f_{\ell_j}^{\hat{v}_j v_j}$ such that $\mathbb{E}[f_{\ell_j}^{\hat{v}_j v_j}] \le (\varepsilon/4)L^{-k}$. Call any expectation of products on the RHS of Equation (63) in which $f_{\ell_j}^{\hat{v}_j v_j}$ occurs as a *light* expectation. Any light expectation is also bounded by $(\varepsilon/4)L^{-k}$. There are at most L^k expectations in the sum. Therefore, losing only an additive factor of $(\varepsilon/4)$ in the upper bound, one can ignore all light expectations on the RHS. The remaining expectations are called *heavy* and are analyzed as follows.

Consider a heavy expectation,

$$\mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}\left[\prod_{j=1}^{k} f_{\ell_j}^{\hat{v}_j v_j}((\overline{r}_j \circ \pi_{\hat{e}_j}))\right] = \mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}\left[\prod_{j=1}^{k}\left(f_{\ell_j}^{\hat{v}_j v_j} \circ \pi_{\hat{e}_j}\right)(\overline{r}_j)\right]. \tag{64}$$

Note that the correlation of the probability space $([L]^k; \tilde{\nu}_{e\rho})$ is at most $\rho < 1$, which depends only on L, k and ε. Thus, applying Lemma 2.9, there is value of γ depending only on L, k and ε, so that,

$$\begin{aligned}\mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}\left[\prod_{j=1}^{k}\left(f_{\ell_j}^{\hat{v}_j v_j} \circ \pi_{\hat{e}_j}\right)(\overline{r}_j)\right] \le \mathbb{E}_{(\overline{r}_1,\ldots,\overline{r}_k)}&\left[\prod_{j=1}^{k} T_{1-\gamma}\left(f_{\ell_j}^{\hat{v}_j v_j} \circ \pi_{\hat{e}_j}\right)(\overline{r}_j)\right]\\ &+ (\varepsilon/4)L^{-k},\end{aligned} \tag{65}$$

where for any $f : [L]^R$ and bijection $\pi : [L] \mapsto [L]$,

$$(f \circ \pi)(\overline{r}) := f(\overline{r} \circ \pi).$$

Note that the ith coordinate of f corresponds to the $\pi(i)$th coordinate of $(f \circ \pi)$. Therefore, Equation (61) implies that for any $1 \le j < j' \le k$,

$$\begin{aligned}\left\{i \;\middle|\; \mathsf{Inf}_i\left(T_{1-\gamma}\left(f_{\ell_j}^{\hat{v}_j v_j} \circ \pi_{\hat{e}_j}\right)\right) > \tau\right\} \bigcap \left\{i \;\middle|\; \mathsf{Inf}_i\left(T_{1-\gamma}\left(f_{\ell_{j'}}^{\hat{v}_{j'} v_{j'}} \circ \pi_{\hat{e}_{j'}}\right)\right) > \tau\right\}&\\ = \emptyset&\end{aligned} \tag{66}$$

Since $(\overline{x}, \overline{y})$ is a δ-smooth solution, the construction of the probability space $([L]^k; \tilde{\nu}_{e\rho})$ implies that measure of its smallest atom is at least $(1-\rho)\left(\delta L^{-1}\right)^k$, which depends only on ε,δ, L and k. Thus, using Equation (66) and setting the value of τ depending only on ε,δ, L and k, one can apply Theorem 2.10 to obtain,

$$\mathbb{E}\left[\prod_{j=1}^{k} T_{1-\gamma}\left(f_{\ell_j}^{\hat{v}_j v_j} \circ \pi_{\hat{e}_j}\right)\right] \le \Gamma_{\overline{\rho}_{k-1}}\left(\mathbb{E}\left[T_{1-\gamma}\left(f_{\ell_1}^{\hat{v}_1 v_1} \circ \pi_{\hat{e}_j}\right)\right], \ldots, \mathbb{E}\left[T_{1-\gamma}\left(f_{\ell_k}^{\hat{v}_k v_k} \circ \pi_{\hat{e}_k}\right)\right]\right) + (\varepsilon/4)L^{-k},$$

where $\overline{\rho}_{k-1} = (\rho, \ldots, \rho)$ is a $(k-1)$-tuple with each entry ρ. Note that the application of the Bonami-Beckner operator does not change the expectation of the above functions, and neither does the permutation of coordinates as each coordinate is sampled u.a.r from the same distribution. Thus,

$$\mathbb{E}\left[T_{1-\gamma}\left(f_{\ell_j}^{\hat{v}_j v_j} \circ \pi_{\hat{e}_j}\right)\right] = \mathbb{E}\left[f_{\ell_j}^{\hat{v}_j v_j}\right], \tag{67}$$

for all $1 \le j \le k$. Therefore, by our assumption, all the expectations in the RHS of the Equation (67) are at least $(\varepsilon/4)L^{-k}$. Applying Lemma 2.4 along with our setting of ρ and using Equation (67) we obtain,

$$\Gamma_{\overline{\rho}_{k-1}}\left(\mathbb{E}\left[T_{1-\gamma}\left(f_{\ell_1}^{\hat{v}_1 v_1} \circ \pi_{\hat{e}_j}\right)\right], \ldots, \mathbb{E}\left[T_{1-\gamma}\left(f_{\ell_k}^{\hat{v}_k v_k} \circ \pi_{\hat{e}_k}\right)\right]\right) \le \left(1 + \frac{1}{k}\right)^{k-1} \prod_{j=1}^{k} \mathbb{E}\left[f_{\ell_j}^{\hat{v}_j v_j}\right]. \tag{68}$$

Combining the above with Equations (67) and (65), we obtain that for the heavy expectations on the RHS of Equation (63),

$$\mathbb{E}\left[\prod_{j=1}^{k} f_{\ell_j}^{\hat{v}_j v_j}\right] \leq \left(1+\frac{1}{k}\right)^{k-1} \cdot \prod_{j=1}^{k} \mathbb{E}\left[f_{\ell_j}^{\hat{v}_j v_j}\right] + (\varepsilon/2)L^{-k}, \leq 3 \cdot \prod_{j=1}^{k} \mathbb{E}\left[f_{\ell_j}^{\hat{v}_j v_j}\right] + (\varepsilon/2)L^{-k}$$

Substituting the above into Equation (63), along with the above observation that the sum of the light expectations is at most $(\varepsilon/4)$, we obtain that the weighted fraction of edges in $E_{\mathcal{H}}$ corresponding to our choice of $\hat{u}$, $(\hat{v}_1, \ldots, \hat{v}_k)$, and $e \in \mathcal{I}$ is satisfied by σ is at most,

$$3 \cdot \sum_{(\ell_1,\ldots,\ell_k)\in C_e} \prod_{j=1}^{k} \mathbb{E}\left[f_{\ell_j}^{\hat{v}_j v_j}\right] + (3\varepsilon/4). \tag{69}$$

Taking the weighted sum of the above over all hyperedges $e \in \mathcal{I}$ completes the proof of the Lemma 5.4. □

For a good vertex $\hat{u} \in V_{\mathcal{U}}$, at least $(1-\sqrt{\eta})$ fraction of k-tuples of its neighbors are good. Therefore, losing an additional additive $\sqrt{\eta}$ in the upper bound, we obtain that the weighted fraction of hyperedges in $E_{\mathcal{H}}$ corresponding to the choice of a good vertex $\hat{u}$ satisfied by σ is at most,

$$\begin{aligned} & 3 \cdot \mathbb{E}_{(\hat{v}_1,\ldots,\hat{v}_k)}\left[\sum_{e=(v_1,\ldots,v_k)\in E_{\mathcal{I}}} w_e \cdot \left[\sum_{(\ell_1,\ldots,\ell_k)\in C_e} \prod_{j=1}^{k} \mathbb{E}\left[f_{\ell_j}^{\hat{v}_j v_j}\right]\right]\right] + \sqrt{\eta} + (3\varepsilon/4) \\ = \; & 3 \cdot \sum_{e=(v_1,\ldots,v_k)\in E_{\mathcal{I}}} w_e \cdot \left[\sum_{(\ell_1,\ldots,\ell_k)\in C_e} \prod_{j=1}^{k} \mathbb{E}_{\hat{v}}\left[\mathbb{E}\left[f_{\ell_j}^{\hat{v} v_j}\right]\right]\right] + \sqrt{\eta} + (3\varepsilon/4), \end{aligned} \tag{70}$$

where $\mathbb{E}_v[.]$ is the expectation over a random neighbor $\hat{v}$ of $\hat{u}$. In the above, the sum over the hyperedges $e \in \mathcal{I}$ is simply the expected number of hyperedges satisfied when each vertex $v \in V_{\mathcal{I}}$ is independently assigned the label ℓ with probability

$$\mathbb{E}_{\hat{v}}\left[\mathbb{E}\left[f_{\ell}^{\hat{v}v}\right]\right].$$

This is at most $\mathsf{opt}(\mathcal{I})$. Moreover, at least $(1-\sqrt{\eta})$ fraction of the vertices $\hat{u}$ are good. Therefore, with an additional loss of $\sqrt{\eta}$ in the upper bound we obtain,

$$\mathsf{val}(\mathcal{H}, \sigma) \leq 3 \cdot \mathsf{opt}(\mathcal{I}) + 2\sqrt{\eta} + (3\varepsilon/4). \tag{71}$$

Choosing ζ to be small enough so that $2\sqrt{\eta} \leq (\varepsilon/4)$ completes the proof of the lemma. □

5.3 Proof of Theorem 3.1

Note that $\mathsf{opt}(\mathcal{I}) \geq L^{-k}$, and since $(\overline{x}, \overline{y})$ is a δ-smooth solution to $\mathsf{LP}(\mathcal{I})$ satisfying Equation (15), one can choose $\delta = 1/2$, and ζ small enough so that Lemma 5.1 implies,

$$\mathsf{opt}(\mathcal{H}) \geq \frac{\rho \cdot \mathsf{lpsup}(\mathcal{I})}{4}, \tag{72}$$

in the YES case.

Also, choosing $\varepsilon = L^{-k}$, Lemma 5.3 implies,

$$\mathsf{opt}(\mathcal{H}) \leq 4 \cdot \mathsf{opt}(\mathcal{I}), \tag{73}$$

in the NO Case. Observing that this setting of ε implies $\rho = \Omega\left(1/\left(k^3 \log L\right)\right)$ proves Theorem 3.1.

6 From a Rounding Gap to an Integrality Gap

Let $\tilde{\mathcal{I}}$ be the instance of CSP-$[\mathcal{C}, k, L]$ consisting of one hyperedge $\tilde{e} = (\tilde{v}_1, \ldots, \tilde{v}_k)$ with a constraint $C_{\tilde{e}}$, and $(\overline{x^*}, \overline{y^*})$ be the solution to $\mathtt{LP}(\tilde{\mathcal{I}})$, as given in Theorem 3.2. This section provides the construction of the integrality gap instance $\mathcal{I}$, followed by the description of the solution $(\overline{x}, \overline{y})$ to $\mathtt{LP}(\mathcal{I})$, and the bound on the optimum of $\mathcal{I}$, as desired in Theorem 3.2.

6.1 Construction of $\mathcal{I}$

For each vertex $\tilde{v}$ of the hyperedge $\tilde{e}$, let $\{S_t^{\tilde{v}} \mid t = 1, \ldots, T\}$ be the corresponding partition of $[L]$ constructed by $\mathsf{Round}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))$. We say that a permutation $\pi : [L] \mapsto [L]$ *respects* the partition $\{S_t\}_{t=1}^T$ if,

$$\ell \in S_t \Leftrightarrow \pi(\ell) \in S_t,$$

for all $\ell \in [L]$ and $t = 1, \ldots, T$. It is easy to see that there are exactly $\prod_{t=1}^r (|S_t|!)$ of such permutations. The following is a randomized construction of $\mathcal{I}$. Here n is a parameter to be set later depending only on L and k.

Vertices. Let $V_j := \{v_{ji} \mid i = 1, \ldots, n\}$, for $j = 1, \ldots, k$ be k layers of vertices. The vertex set is their union, i.e., $V_{\mathcal{I}} = \cup_{j=1}^k V_j$.

Hyperedges. For every $(i_1, \ldots, i_k) \in [n]^k$ there is a hyperedge $e = (v_{1i_1}, \ldots, v_{ki_k})$. The constraint C_e is chosen independently at random as follows. Choose a $\{S_t^{\tilde{v}_j}\}_{t=1}^T$ respecting permutation π_j uniformly at random, and independently for $j = 1, \ldots, k$, and let,

$$C_e = [\pi_1, \ldots, \pi_k] C_{\tilde{e}}. \tag{74}$$

Assign to each of the n^k hyperedges in $\mathcal{I}$ the same weight n^{-k}.

6.2 LP Solution for $\mathcal{I}$

Let us first create $(\overline{\tilde{x}}, \overline{\tilde{y}})$ as solution to to the relaxation $\mathtt{LP}_1(\tilde{\mathcal{I}})$, given in Section 7. Let $(\overline{\hat{x}}, \overline{\hat{y}})$ be the 0.1-smooth solution constructed in Step 1 of $\mathsf{Round}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))$. For each $\overline{\ell} \in [L]^k$ let,

$$\tilde{y}_{\tilde{e}\overline{\ell}} = \frac{\hat{y}_{\tilde{e}\overline{\ell}}}{2}. \tag{75}$$

For each vertex $\tilde{v}_j$ $(1 \le j \le k)$ in $\tilde{\mathcal{I}}$, and $\ell \in [L]$ such that $\ell \in S_t^{\tilde{v}}$, let,

$$\tilde{x}_{\tilde{v}_j \ell} = \left(1/2^t\right). \tag{76}$$

Observe that $\tilde{x}_{\tilde{v}_j \ell} \ge (1/2)\hat{x}_{\tilde{v}_j \ell}$. Along with Equation (75) this implies that $(\overline{\tilde{x}}, \overline{\tilde{y}})$ is a valid solution to $\mathtt{LP}_1(\tilde{\mathcal{I}})$. Furthermore,

$$\mathsf{lpval}_1(\tilde{\mathcal{I}}, (\overline{\tilde{x}}, \overline{\tilde{y}})) = \frac{\mathsf{lpval}(\tilde{\mathcal{I}}, (\overline{\hat{x}}, \overline{\hat{y}})}{2}, \tag{77}$$

where lpval_1 is as defined in Section 7. A solution $(\overline{x'}, \overline{y'})$ to the relaxation $\mathtt{LP}_1(\mathcal{I})$ is constructed as follows. Let $e = (v_1, \ldots, v_k)$ be a hyperedge in $\mathcal{I}$, where $v_j \in V_j$ $(1 \le j \le k)$. The corresponding constraint

C_e is given by $[\pi_1, \ldots, \pi_k]C_{\tilde{e}}$ where π_j respects the partition $\{S_t^{\tilde{v}_j}\}_{t=1}^T$, for $j = 1, \ldots, k$. For each $\overline{\ell} = (\ell_1, \ldots, \ell_k)$ let $\overline{\ell}' = (\pi_1^{-1}(\ell_1), \ldots, \pi_k^{-1}(\ell_k))$, so that,

$$\overline{\ell}' \in C_{\tilde{e}} \Leftrightarrow \overline{\ell} \in [\pi_1, \ldots, \pi_k]C_{\tilde{e}} = C_e, \tag{78}$$

and let,

$$y'_{e\overline{\ell}} = \tilde{y}_{\tilde{e}\overline{\ell}'}. \tag{79}$$

Essentially, the LP variables corresponding to the hyperedges are permuted according to the sequence of permutations used in constructing the hyperedge. On the other hand, since the permutations π_j respects $\{S_t^{\tilde{v}_j}\}_{t=1}^T$ $(1 \le j \le k)$, the variables corresponding to the vertices do not change. Formally, for each $v \in V_j$ $(1 \le j \le k)$ and $\ell \in [L]$,

$$x'_{v\ell} = \tilde{x}_{\tilde{v}_j\ell}. \tag{80}$$

Note that for a given $t \in \{1, \ldots, T\}$, for all $\ell \in S_t^{\tilde{v}_j}$, $x'_{v\ell}$ has the same value. Along with the fact that the permutations π_j used to construct C_e respect $\{S_t^{\tilde{v}_j}\}_{t=1}^T$ $(1 \le j \le k)$, this implies that $(\overline{x'}, \overline{y'})$ is a valid solution to $\mathtt{LP}_1(\mathcal{I})$. From the construction of C_e we have,

$$\sum_{\overline{\ell} \in C_e} y'_{e\overline{\ell}} = \sum_{\overline{\ell}' \in C_{\tilde{e}}} \tilde{y}_{\tilde{e}\overline{\ell}'} = \mathsf{lpval}_1(\tilde{\mathcal{I}}, (\overline{\tilde{x}}, \overline{\tilde{y}})). \tag{81}$$

Since each hyperedge in $\mathcal{I}$ has the same normalized weight, we obtain,

$$\begin{aligned} \mathsf{lpval}_1(\mathcal{I}, (\overline{x'}, \overline{y'})) = \mathsf{lpval}_1(\tilde{\mathcal{I}}, (\overline{\tilde{x}}, \overline{\tilde{y}})) = \frac{\mathsf{lpval}(\tilde{\mathcal{I}}, (\overline{\hat{x}}, \overline{\hat{y}}))}{2} \quad &\ge \quad \frac{0.9\mathsf{lpval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))}{2} \\ &\ge \quad \frac{\mathsf{lpval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))}{4}, \end{aligned} \tag{82}$$

where the second last inequality follows from the fact that $(\overline{\hat{x}}, \overline{\hat{y}})$ is 0.1-smooth solution corresponding to $(\overline{x^*}, \overline{y^*})$ and Lemma 2.2. Applying Lemma 7.1 to Equation (82) yields a solution $(\overline{x}, \overline{y})$ to $\mathtt{LP}(\mathcal{I})$ such that,

$$\mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})) \ge \frac{\mathsf{lpval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))}{4}. \tag{83}$$

6.3 Bound on $\mathsf{opt}(\mathcal{I})$

Consider a fixed labeling $\sigma : V_{\mathcal{I}} \mapsto [L]$. We shall estimate the number of hyperedges in $\mathcal{I}$ satisfied by σ over the random choice of the constraints as given in the construction of $\mathcal{I}$, and show that this does not deviate much from the expectation, except with very low probability. A further application of union-bound yields the desired upper bound.

Let $e = (v_1, \ldots, v_k) \in E_{\mathcal{I}}$, where $v_j \in V_j$ for $j = 1, \ldots, k$. Let $t_j \in \{1, \ldots, T\}$ be such that $\sigma(v_j) \in S_{t_j}^{\tilde{v}_j}$ for $j = 1, \ldots, k$. Let p_e be the probability over the choice of C_e that σ satisfies e.

Lemma 6.1. *Either $p_e = 0$ or $p_e \ge L^{-k}$.*

Proof. It is easy to see that,

$$\begin{aligned} \left(\prod_{j=1}^{k} S_{t_j}^{\tilde{v}_j}\right) \bigcap C_{\tilde{e}} = \emptyset \quad \Leftrightarrow \quad & \forall (\pi_1, \ldots \pi_k) \text{ s.t. } \pi_j \text{ respects } \{S_t^{\tilde{v}_j}\}_{t=1}^T, j = 1, \ldots, k, \\ & \sigma \text{ does not satisfy } [\pi_1, \ldots, \pi_k]C_{\tilde{e}}. \end{aligned} \tag{84}$$

Thus, if the LHS of Equation (84) holds for e, then $p_e = 0$. Otherwise, with probability at least,

$$\prod_{j=1}^{k} \left| S_{t_j}^{\tilde{v}_j} \right|^{-1} \geq L^{-k},$$

over the choice of $\{\pi_j\}_{j=1}^{k}$, $(\sigma(v_1), \ldots, \sigma(v_k)) \in C_e$. $\square$

We also have the following lemma.

Lemma 6.2.

$$p_e \leq T^k \cdot \mathsf{Roundval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*})). \tag{85}$$

Proof. If $p_e = 0$ then the lemma is trivially true. Otherwise, from the Equation (84), $S_{t_j}^{\tilde{v}_j} \neq \emptyset$, for $j = 1, \ldots, k$. By the randomized construction of C_e, it can be seen that p_e is the probability that $(\ell_1, \ldots, \ell_k) \in C_{\tilde{e}}$, when ℓ_j is chosen independently and u.a.r from $S_{t_j}^{\tilde{v}_j}$. This is same as the probability that the algorithm $\mathsf{Round}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))$ satisfies $\tilde{e}$, after choosing the index t_j for $\tilde{v}_j$ in Step 2b for $j = 1, \ldots, k$. Since this choice is made with probability at least T^{-k}, the lemma follows. $\square$

The following key lemma gives the desired bound on the probability that the number hyperedges satisfied is much larger than expected.

Lemma 6.3. *For any $\varepsilon \in (0, 1)$, there is a value of n depending only on L, k, and ε, such that,*

$$\Pr\left[\text{Weighted fraction of hyperedges in } \mathcal{I} \text{ satisfied by } \sigma > (1+\varepsilon)\mathcal{R}\right] < L^{-kn},$$

where $\mathcal{R} := T^k \cdot \mathsf{Roundval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*}))$, and the probability is taken over the choice of the constraints C_e, $e \in E_{\mathcal{I}}$.

Proof. We may assume that,

$$|\{e \in E_{\mathcal{I}} \mid p_e > 0\}| \geq n^k \cdot \mathcal{R}, \tag{86}$$

otherwise the lemma follows trivially as each edge has weight n^{-k}. Since $C_{\tilde{e}} \neq \emptyset$, it can be seen from the description of Round in Figure 2 that,

$$\mathsf{Roundval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*})) \geq T^{-k} L^{-k}, \tag{87}$$

which along with Equation (86), Lemma 6.1, and the setting of $\mathcal{R}$ implies,

$$\sum_{e \in E_{\mathcal{I}}} p_e \geq n^k L^{-2k}. \tag{88}$$

Observe that the choice of C_e and therefore the event that e is satisfied by σ is independent for all hyperedges. Therefore, applying the Chernoff bound we have,

$$\Pr\left[|\{e \mid e \text{ satisfied by } \sigma\}| > (1+\varepsilon) \sum_{e \in E_{\mathcal{I}}} p_e\right] < \exp\left(\frac{-\varepsilon^2 \cdot \sum_{e \in E_{\mathcal{I}}} p_e}{3}\right). \tag{89}$$

Choosing n large enough depending only on L, k and ε and substituting in the above from Equation (88) completes the proof of the lemma. $\square$

Let us fix $\varepsilon = 1/2$. Note that from the description of Round in Figure 2, $T = O(\log L)$. Observing that the number of vertices in $\mathcal{I}$ is nk and the total number of labelings of its vertices is L^{kn}, applying the union bound to Lemma 6.3 yields the bound on $\mathsf{opt}(\mathcal{I})$.

Lemma 6.4. *For a large enough value of n depending only on L and k, there exists an instance $\mathcal{I}$ whose constraints are permutations of $C_{\tilde{e}}$ such that,*

$$\mathsf{opt}(\mathcal{I}) = O\left((\log L)^k\right) \mathsf{Roundval}(\tilde{\mathcal{I}}, (\overline{x^*}, \overline{y^*})). \tag{90}$$

7 Relaxation $\mathtt{LP}_1$

Figure 3 gives an alternate LP relaxation, $\mathtt{LP}_1$ for CSP-$[\mathcal{C}, k, L]$, in which the constraints with equality in $\mathtt{LP}$ are further relaxed. Let $\mathsf{lpval}_1(\mathcal{I}, (\overline{x}, \overline{y}))$ denote the objective value of $\mathtt{LP}_1(\mathcal{I})$ on the solution $(\overline{x}, \overline{y})$,

$$\max \sum_{e \in E_{\mathcal{I}}} w_e \cdot \sum_{\overline{\ell} \in C_e} y_{e\overline{\ell}} \tag{91}$$

subject to,

$$\forall v \in V_{\mathcal{I}}, \qquad \sum_{\ell \in [L]} x_{v\ell} \le 1 \tag{92}$$

$$\begin{aligned} &\forall v \in V_{\mathcal{I}} \text{ and,} \\ &e = (v_1, \ldots, v_{i-1}, v, v_{i+1}, \ldots, v_k) \in E_{\mathcal{I}} \text{ and,} \\ &\ell^* \in [L], \qquad \sum_{\overline{\ell} \in [L]^{i-1} \times \{\ell^*\} \times [L]^{k-i}} y_{e\overline{\ell}} \le x_{v\ell^*} \end{aligned} \tag{93}$$

$$\forall v \in V_{\mathcal{I}}, \ell \in [L], \qquad x_{v\ell} \ge 0. \tag{94}$$

$$\forall e \in E_{\mathcal{I}}, \overline{\ell} \in [L]^k, \qquad y_{e\overline{\ell}} \ge 0. \tag{95}$$

Figure 3: LP Relaxation $\mathtt{LP}_1(\mathcal{I})$ for instance $\mathcal{I}$ of CSP-$[\mathcal{C}, k, L]$.

and $\mathsf{lpsup}_1(\mathcal{I})$ its supremum over all $(\overline{x}, \overline{y})$. The following lemma states that with regards to the optimum objective value, $\mathtt{LP}$ and $\mathtt{LP}_1$ are equivalent.

Lemma 7.1. *For any instance $\mathcal{I}$ of* CSP-$[\mathcal{C}, k, L]$*, if $(\overline{x'}, \overline{y'})$ is a solution to $\mathtt{LP}_1(\mathcal{I})$, then there exists a solution $(\overline{x}, \overline{y})$ to $\mathtt{LP}(\mathcal{I})$ such that,*

$$\mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})) \ge \mathsf{lpval}_1(\mathcal{I}, (\overline{x'}, \overline{y'})). \tag{96}$$

In particular,

$$\mathsf{lpsup}_1(\mathcal{I}) = \mathsf{lpsup}(\mathcal{I}).$$

Proof. Let $(\overline{x'}, \overline{y'})$ be as given in the statement of the lemma. We can make tight all the constraints given by Equation (92) by choosing some $\ell \in [L]$ and if needed increase $x_{v\ell}$ so that $\sum_{\ell \in [L]} x'_{v\ell} = 1$ for each $v \in V_{\mathcal{I}}$. Now, let $e = (v_1, \dots, v_k)$, $t \in [k]$ and $\ell_t^* \in [L]$ such that,

$$\sum_{\overline{\ell} \in [L]^{t-1} \times \{\ell_t^*\} \times [L]^{k-t}} y'_{e\overline{\ell}} < x'_{v_t \ell_t^*}$$

The above implies that,

$$\sum_{\overline{\ell} \in [L]^k} y'_{e\overline{\ell}} < \sum_{\ell \in [L]} x'_{v_t \ell}$$

Since the RHS of the above equals 1 for each $v_1, \dots, v_k$, this further implies that for each $i \in [k]$ there is $\ell_i^* \in [L]$ such that,

$$\sum_{\overline{\ell} \in [L]^{i-1} \times \{\ell_i^*\} \times [L]^{k-i}} y'_{e\overline{\ell}} < x'_{v_t \ell_i^*} \tag{97}$$

Let $\overline{\ell^*} = (\ell_1^*, \dots, \ell_k^*)$. The variable $y'_{e\overline{\ell^*}}$ can be increase so that (97) becomes tight for at least one $i \in [k]$.

The above procedure can continue by increasing the $\{y'_{e\ell}\}$ variables till all the constraints given by Equation (93) become tight in which case we obtain a solution for the relaxation $\texttt{LP}(\mathcal{I})$. Since the variables are only increased this preserves the objective value. □

8 Proof of Theorem 2.10

We shall require the following bi-linear Gaussian stability bound as shown in [14] (as Theorem 1.14 and Proposition 1.15).

Theorem 8.1. *Let $(\Omega_i^{(1)} \times \Omega_i^{(2)}, \mu_i)$ be a sequence of correlated spaces such that for each i, the probability of any atom in $(\Omega_i^{(1)} \times \Omega_i^{(2)}, \mu_i)$ is at least $\alpha \leq 1/2$ and such that $\rho(\Omega_i^{(1)}, \Omega_i^{(2)}; \mu_i) \leq \rho$ for all i. Then there exists a universal constant C such that, for every $\nu > 0$, taking*

$$\tau = \nu^{\left(C \frac{\log(1/\alpha)\log(1/\nu)}{\nu(1-\rho)}\right)},$$

for functions $f : \prod_{i=1}^n \Omega_i^{(1)} \mapsto [0,1]$ and $g : \prod_{i=1}^n \Omega_i^{(2)} \mapsto [0,1]$ that satisfy,

$$\min(\mathsf{Inf}_i(f), \mathsf{Inf}_i(g)) \leq \tau,$$

for all i, we have,

$$\mathbb{E}[fg] \leq \Gamma_\rho(\mathbb{E}[f], \mathbb{E}[g]) + \nu.$$

The proof of Theorem 2.10 uses the following lemma on the influences of a product of functions, proved in [14] (as Lemma 6.5).

Lemma 8.2. *Let $f_1, \dots, f_k : \Omega^n \mapsto [0,1]$. Then for all $i = 1, \dots, n$:*

$$\mathsf{Inf}_i\left(\prod_{j=1}^k f_j\right) \leq k \sum_{j=1}^k \mathsf{Inf}_i(f_j). \tag{98}$$

Define for each $j = 1, \ldots, k-1$,

$$f_{>j} := \prod_{j'=j+1}^{k} f_{j'}. \tag{99}$$

We have the following lemma.

Lemma 8.3. *For all $j = 1, \ldots, k-1$,*

$$\min(\mathsf{Inf}_i(f_j), \mathsf{Inf}_i(f_{>j})) \le k^2\tau, \tag{100}$$

for any $i = 1, \ldots, n$.

Proof. Suppose $\mathsf{Inf}_i(f_j) > k^2\tau$. Then, Equation (22) implies that $\mathsf{Inf}_i(f_{j'}) < \tau$ for all $j' = j+1, \ldots, k$. Using Lemma 8.2 along with the definition of $f_{>j}$ yields $\mathsf{Inf}_i(f_{>j}) \le k^2\tau$.

On the other hand, if $\mathsf{Inf}_i(f_{>j}) > k^2\tau$, then – again by Lemma 8.2 – there must be some $j' \in \{j+1, \ldots, k\}$ such that $\mathsf{Inf}_i(f_{j'}) > \tau$, and thus Equation (22) implies $\mathsf{Inf}_i(f_j) \le \tau$. □

With the setting of τ as given in (21), recursively applying Theorem 8.1 to $\mathbb{E}[f_{>j}]$ for $j = 1, \ldots, k-1$ we obtain,

$$\begin{aligned}
\mathbb{E}\left[\prod_{j=1}^{k} f_j\right] &= \mathbb{E}\left[f_1 f_{>1}\right] \\
&\le \Gamma_\rho(\mathbb{E}[f_1], \mathbb{E}[f_{>1}]) + (\nu/k) \\
&\le \Gamma_\rho(\mathbb{E}[f_1], \Gamma_\rho(\mathbb{E}[f_2], \mathbb{E}[f_{>2}]) + (\nu/k)) + (\nu/k) \\
&\le \Gamma_\rho(\mathbb{E}[f_1], \Gamma_\rho(\mathbb{E}[f_2], \mathbb{E}[f_{>2}])) + (2\nu/k) \\
&\le \vdots \\
&\le \Gamma_{\rho,\ldots,\rho}(\mathbb{E}[f_1], \ldots, \mathbb{E}[f_k]) + \nu,
\end{aligned} \tag{101}$$

where the last inequality is obtained by collecting the $(k-1)$ error terms outside which sum up to $((k-1)\nu/k) \le \nu$.

9 Proof of Lemma 2.4

Let $\psi(t) := (1/\sqrt{2\pi})e^{-t^2/2}$ denote the probability density function of a standard Gaussian random variable; $\Phi(t)$ be its cumulative distribution function and let $\tilde{\Phi}(t)$ be the probability that a standard Gaussian random variable is at least t, i.e. $\tilde{\Phi}(t) = 1 - \Phi(t) = \Phi(-t)$. The following lemma (proved as Lemma A.1 in [2]) shows useful bounds on these functions.

Lemma 9.1. *For every $t > 0$*

$$\frac{t \cdot \psi(t)}{t^2+1} < \tilde{\Phi}(t) < \frac{\psi(t)}{t}; \tag{102}$$

and therefore, for every $p \ge 2$,

$$\tilde{\Phi}^{-1}(1/p) \le c\sqrt{\log(p)}, \tag{103}$$

for some universal constant $c > 0$.

For our analysis we shall need bounds for the Gaussian stability $\Gamma_\rho(\mu, \nu)$ (see Definition 2.3). Note that since $\rho \in [0, 1]$, $\Gamma_\rho(\mu, \nu) \geq \mu\nu$. The following lemma shows that the Gaussian random variables in Equation (17) can be truncated while essentially preserving the LHS.

Lemma 9.2. *Let $T \geq 2$ and $\mu, \nu \geq 1/T$. Then,*
(i) $\Phi^{-1}(\mu), \Phi^{-1}(\nu) \geq -c\sqrt{\log T}$.
(ii) *Fix any $\delta \in (0, 1]$ and let,*

$$\kappa := c\sqrt{2\log T + \log(3/\delta)}, \tag{104}$$

$$a := \min\{\Phi^{-1}(\mu), \kappa\} \text{ and, } \ b := \min\{\Phi^{-1}(\nu), \kappa\}. \tag{105}$$

Then,

$$\Pr[-\kappa \leq X \leq a, Y \leq b] \geq (1 - \delta)\Gamma_\rho(\mu, \nu),$$

where X and Y are standard Gaussian random variables with correlation $\rho \in [0, 1]$. Here, c is the constant from Lemma 9.1.

Proof. (i) From Equation (103) of Lemma 9.1, we have $\tilde{\Phi}^{-1}(1/T) \leq c\sqrt{\log T}$. Since $\mu \geq 1/T$, $\Phi^{-1}(\mu) \geq \Phi^{-1}(1/T) = -\tilde{\Phi}^{-1}(1/T)$. Thus, $\Phi^{-1}(\mu) \geq -c\sqrt{\log T}$, and similarly for ν.

(ii) From Equation (103) of Lemma 9.1,

$$\tilde{\Phi}\left(\kappa\right) \leq \frac{\delta}{3T^2}.$$

Observe that,

$$\begin{aligned}
\Gamma_\rho(\mu, \nu) &= \Pr[X \leq \Phi^{-1}(\mu), Y \leq \Phi^{-1}(\nu)], \\
&\leq \Pr[-\kappa \leq X \leq a, Y \leq b] + \Pr[X < -\kappa] + \Pr[X > \kappa] + \Pr[Y > \kappa], \\
&= \Pr[-\kappa \leq X \leq a, Y \leq b] + 3\tilde{\Phi}\left(\kappa\right), \\
&\leq \Pr[-\kappa \leq X \leq a, Y \leq b] + \frac{\delta}{T^2}, \qquad (106)
\end{aligned}$$

and that,

$$\Gamma_\rho(\mu, \nu) \geq \mu\nu \geq \frac{1}{T^2},$$

which completes the proof of the lemma. $\square$

Using the above lemma, we prove the following key upper bound on Gaussian stability.

Lemma 9.3. *Let $T \geq 2$ and $1 \geq \mu, \nu \geq (1/T)$. There is a universal constant $C > 0$ such that, for any $\varepsilon \in (0, 1/2]$,*

$$\rho = \frac{\varepsilon}{C(\log T + \log(1/\varepsilon))}, \tag{107}$$

implies,

$$\Gamma_\rho(\mu, \nu) \leq (1 + \varepsilon)\mu\nu.$$

Proof. Applying Lemma 9.2 shows that letting,

$$\kappa = c\sqrt{2\log T + \log(12/\varepsilon)}, \tag{108}$$

and the corresponding values of a and b as given in Equation (105), yields,

$$\Pr[-\kappa \le X \le a, Y \le b] \ge (1-\varepsilon/4)\Gamma_\rho(\mu,\nu), \tag{109}$$

where X and Y are standard Gaussian random variables with correlation $\rho \in [0,1]$. We have the following lemma (proved below).

Lemma 9.4. *Setting ρ as given in Equation* (107) *implies,*

$$\Pr[-\kappa \le X \le a, Y \le b] \le (1+\varepsilon/2)\mu\nu. \tag{110}$$

Combining Equations (109) and (110) we obtain,

$$\Gamma_\rho(\mu,\nu) \;\le\; \frac{(1+\varepsilon/2)}{(1-\varepsilon/4)}\mu\nu \;\le\; (1+\varepsilon)\mu\nu, \tag{111}$$

using the fact that $(1+\varepsilon/2) \le (1-\varepsilon/4)(1+\varepsilon)$ for $\varepsilon \in (0,1/2]$, thus completing the proof of Lemma 9.3. $\square$

Proof. (of Lemma 9.4) Since X and Y are ρ-correlated, $Y = \rho X + \sqrt{1-\rho^2}Z$, where Z is a standard Gaussian random variable independent of X. Thus,

$$Y \le b \;\Leftrightarrow\; \rho X + \sqrt{1-\rho^2}Z \le b \;\Leftrightarrow\; Z \le \frac{b-\rho X}{\sqrt{1-\rho^2}}. \tag{112}$$

Therefore,

$$\begin{aligned}
&\Pr[-\kappa \le X \le a, Y \le b] \\
=\;& \Pr\left[-\kappa \le X \le a, Z \le \frac{b-\rho X}{\sqrt{1-\rho^2}}\right], \\
\le\;& \Pr\left[-\kappa \le X \le a, Z \le \frac{b+\rho \kappa}{\sqrt{1-\rho^2}}\right] \quad \text{(since } |X| \le \kappa\text{)}, \\
=\;& \Pr\left[-\kappa \le X \le a\right]\Pr\left[Z \le \frac{b+\rho \kappa}{\sqrt{1-\rho^2}}\right].
\end{aligned} \tag{113}$$

Observing that $\Pr\left[-\kappa \le X \le a\right] \le \mu$ and $\Pr\left[Z \le b\right] \le \nu$, application of Lemma 9.5 proved below completes the proof of Lemma 9.4. $\square$

Lemma 9.5. *For the above setting of parameters the following holds.*

$$\Pr\left[Z \le \frac{b+\rho\kappa}{\sqrt{1-\rho^2}}\right] \le \left(1+\frac{\varepsilon}{2}\right)\Pr\left[Z \le b\right].$$

Proof. For convenience let,

$$b' = \frac{b+\rho\kappa}{\sqrt{1-\rho^2}},$$

which implies,

$$\begin{aligned}
|b' - b| &= \left|\frac{b+\rho\kappa}{\sqrt{1-\rho^2}} - b\right|, \\
&= \left|\frac{b+\rho\kappa - b\sqrt{1-\rho^2}}{\sqrt{1-\rho^2}}\right|, \\
&\leq \frac{\left|b - b\sqrt{1-\rho^2}\right| + \rho\kappa}{\sqrt{1-\rho^2}} \\
&\leq \frac{|b|(1-\sqrt{1-\rho^2}) + \rho\kappa}{\sqrt{1-\rho^2}} \\
&\leq \frac{|b|(1-(1-\rho^2)) + \rho\kappa}{\sqrt{1-\rho^2}} \\
&= \frac{|b|\rho^2 + \rho\kappa}{\sqrt{1-\rho^2}} \qquad (114)
\end{aligned}$$

We consider the following two cases.

Case 1: $|b| < 10$. This implies that,

$$\Pr\left[Z \leq b\right] \geq c^*, \qquad (115)$$

where c^* is an absolute constant. On the other hand observe that $|b| \leq \kappa$ and thus, ρ, $\rho\kappa$ and $|\rho^2 b|$ can be made small enough by the choice of the constant C in Lemma 9.3 so that,

$$\left|\Pr\left[Z \leq b'\right] - \Pr\left[Z \leq b\right]\right| \leq \left|b' - b\right| \leq (\varepsilon c^*)/2. \qquad (116)$$

Combining Equations (115) and (116) proves the lemma for this case.

Case 2: $|b| \geq 10$. In this case, using Equation (114), choosing the constant C to be large enough we can ensure that,

$$\text{sign}(b) = \text{sign}(b').$$

In particular, the above implies that,

$$b^* := \arg\max_{x \in [b, b']} \psi(x) \Rightarrow b^* \in \{b, b'\}. \qquad (117)$$

Thus,

$$\left|\Pr\left[Z \leq b'\right] - \Pr\left[Z \leq b\right]\right| = \left|\Phi(b') - \Phi(b)\right| \leq |b' - b|\psi(b^*). \qquad (118)$$

Diving the above by $\Phi(b)$ we obtain,

$$\begin{aligned}
\left|\frac{\Phi(b') - \Phi(b)}{\Phi(b)}\right| &\leq \frac{|b' - b|\psi(b^*)}{\Phi(b)} \\
&\leq \frac{|b' - b|\psi(b^*)}{\tilde{\Phi}(|b|)} \\
&\leq \frac{(b^2+1)|b' - b|\psi(b^*)}{|b| \cdot \psi(|b|)}, \qquad (119)
\end{aligned}$$

where the second last inequality follows from the fact that $\Phi(b) \geq \tilde{\Phi}(|b|)$ for $|b| > 0$, and the last inequality follows from the bound lower bound in Equation (102). Note that,

$$|b'^2 - b^2| = \left| \left(\frac{b + \rho\kappa}{\sqrt{1-\rho^2}} \right)^2 - b^2 \right| \leq \frac{|b\rho|^2 + 2|b\kappa\rho| + |\rho\kappa|^2}{1-\rho^2}. \tag{120}$$

Since $|b| \leq \kappa$, by our setting of ρ in Equation (107) and κ in Equation (108), choosing a large enough value of C we ensure that the RHS of Equation (120) is at most $1/4$. From the definition of b^*, this implies,

$$\frac{\psi(b^*)}{\psi(|b|)} \leq e^{1/8} \leq 5/4. \tag{121}$$

Further, for $|b| > 10$, from Equation (114), we have,

$$\frac{(b^2+1)|b'-b|}{|b|} \leq 2|b||b'-b| \leq \frac{2|b|^2\rho^2 + 2|b|\rho\kappa}{\sqrt{1-\rho^2}} \tag{122}$$

Observe that by a large enough choice of C, both $|b|^2\rho^2$ and $|b|\rho\kappa$ can be bounded from the above by $\varepsilon/20$, and $\sqrt{1-\rho^2}$ be made least $4/5$ yielding,

$$\frac{(b^2+1)|b'-b|}{|b|} \leq \frac{\varepsilon}{4}. \tag{123}$$

Combining the above with Equations (121) and (119) gives us,

$$\left|\Phi(b') - \Phi(b)\right| \leq \left(\frac{\varepsilon}{2}\right)\Phi(b), \tag{124}$$

which completes the proof of the lemma. □

We are ready to prove Lemma 2.4 which is restated as follows.

Lemma 9.6. *Let $k \geq 2$, $T \geq 2$, and $1 \geq \mu_i \geq (1/T)$ for $i = 1, \ldots, k$. Then for any $\varepsilon \in (0, 1/2]$, setting,*

$$\rho = \frac{\varepsilon}{(k-1)C(\log T + \log(1/\varepsilon))}, \tag{125}$$

implies,

$$\Gamma_{\overline{\rho}_{k-1}}(\mu_1, \ldots, \mu_k) \leq (1+\varepsilon)^{k-1} \prod_{i=1}^{k} \mu_i,$$

where $\overline{\rho}_{k-1} = (\rho, \ldots, \rho)$, is a $(k-1)$-tuple with each entry ρ. In Equation (125)*, C is the constant from Lemma 9.3.*

Proof. The proof proceeds via induction on k. For $k = 2$, Lemma 9.3 yields the proof. Assume that the lemma holds for $(k-1) \geq 2$. For k, we have by definition (Equation (18)),

$$\Gamma_{\overline{\rho}_{k-1}}(\mu_1, \ldots, \mu_k) = \Gamma_\rho(\mu_1, \Gamma_{\overline{\rho}_{k-2}}(\mu_2, \ldots, \mu_k)). \tag{126}$$

Let us define,

$$\rho' := \frac{\varepsilon}{(k-2)C(\log T + \log(1/\varepsilon))}.$$

Since $0 \leq \rho < \rho'$, from the inductive definition in Equation (18) and an application of Lemma 9.7 it is easy to see that,

$$\Gamma_{\overline{\rho}_{k-2}}(\mu_2, \ldots, \mu_k) \leq \Gamma_{\overline{\rho'}_{k-2}}(\mu_2, \ldots, \mu_k), \tag{127}$$

where $\overline{\rho'}^{k-2}$ is a $(k-2)$-tuple with each entry ρ'. Applying the inductive hypothesis for $(k-2)$ we obtain,

$$\Gamma_{\overline{\rho'}_{k-2}}(\mu_2, \ldots, \mu_k) \leq (1+\varepsilon)^{k-2} \prod_{i=2}^{k} \mu_i,$$

which in conjunction with Equation (127) gives us,

$$\Gamma_{\overline{\rho}_{k-2}}(\mu_2, \ldots, \mu_k) \leq (1+\varepsilon)^{k-2} \prod_{i=2}^{k} \mu_i. \tag{128}$$

Since $\rho \geq 0$, it is easy to see that,

$$\Gamma_{\overline{\rho}_{k-2}}(\mu_2, \ldots, \mu_k) \geq \prod_{i=2}^{k} \mu_i \geq \left(\frac{1}{T}\right)^{k-1}. \tag{129}$$

Further $\mu_1 \geq (1/T)$, and applying the Lemma 9.3 to the RHS of Equation (126), we obtain,

$$\begin{aligned}
\Gamma_{\overline{\rho}_{k-1}}(\mu_1, \ldots, \mu_k) &\leq (1+\varepsilon)\mu_1 \Gamma_{\overline{\rho}_{k-2}}(\mu_2, \ldots, \mu_k) \\
&\leq (1+\varepsilon)\mu_1 \cdot (1+\varepsilon)^{k-2} \prod_{i=2}^{k} \mu_i \quad \text{By Equation (128),} \\
&= (1+\varepsilon)^{k-1} \prod_{i=1}^{k} \mu_i, \qquad (130)
\end{aligned}$$

which completes the inductive step. $\square$

Lemma 9.7. *For $\mu, \nu \in [0,1]$, and $1 \leq \rho < \rho' \leq 1$,*

$$\Gamma_\rho(\mu, \nu) \leq \Gamma_{\rho'}(\mu, \nu).$$

Proof. (Sketch) The lemma is obtained by differentiating $\Gamma_\rho(\mu, \nu)$ with respect to ρ and showing that it is non-negative in the range $[0, 1)$. We omit the details. $\square$

10 Proof of Lemma 2.2

Proof. Given a solution $(\overline{x^*}, \overline{y^*})$, construct a valid δ-smooth solution $(\overline{x}, \overline{y})$ as follows:

$$y_{e\overline{\ell}} = (1-\delta) y^*_{e\overline{\ell}} + \delta L^{-k}, \quad \forall e \in E, \overline{\ell} \in [L]^k, \tag{131}$$

and,

$$x_{v\ell} = (1-\delta) x^*_{v\ell} + \delta L^{-1}, \quad \forall v \in V, \ell \in [L]. \tag{132}$$

It is easy to see that the objective value decreases by at most a multiplicative factor of δ. By the definition of $\mathsf{lpval}(\mathcal{I})$ and since the set of all valid solutions to $\mathtt{LP}(\mathcal{I})$ is a closed set, there must be a solution $(\overline{x^*}, \overline{y^*})$ such that,

$$\mathsf{lpval}(\mathcal{I}, (\overline{x^*}, \overline{y^*})) = \mathsf{lpsup}(\mathcal{I}). \tag{133}$$

We use $(\overline{x^*}, \overline{y^*})$ to construct $(\overline{x}, \overline{y})$ as above which yields,

$$\mathsf{lpval}(\mathcal{I}, (\overline{x}, \overline{y})) \geq (1 - \delta)\mathsf{lpsup}(\mathcal{I}),$$

proving Equation (15). □

Acknowledgment

The authors thank Elchanan Mossel for helpful discussion on Gaussian stability bounds.

References

[1] S. O. Chan. Approximation resistance from pairwise independent subgroups. In *Proc. STOC*, pages 447–456, 2013.

[2] M. Charikar, K. Makarychev, and Y. Makarychev. Near-optimal algorithms for unique games. In *Proc. STOC*, pages 205–214, 2006.

[3] V. Dalmau and A. A. Krokhin. Robust satisfiability for CSPs: Hardness and algorithmic results. *TOCT*, 5(4):15, 2013.

[4] V. Dalmau, A. A. Krokhin, and R. Manokaran. Towards a characterization of constant-factor approximable min CSPs. In *Proc. SODA*, pages 847–857, 2015.

[5] U. Feige and G. Schechtman. On the optimality of the random hyperplane rounding technique for MAX CUT. *Random Struct. Algorithms*, 20(3):403–440, 2002.

[6] M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM*, 42(6):1115–1145, 1995.

[7] S. Khot. On the power of unique 2-prover 1-round games. In *Proc. STOC*, pages 767–775, 2002.

[8] S. Khot, G. Kindler, E. Mossel, and R. O'Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM Journal of Computing*, 37(1):319–357, 2007.

[9] S. Khot and N. K. Vishnoi. The unique games conjecture, integrality gap for cut problems and embeddability of negative type metrics into ℓ_1. In *Proc. FOCS*, pages 53–62, 2005.

[10] G. Kindler, A. Kolla, and L. Trevisan. Approximation of non-boolean 2CSP. *CoRR*, abs/1504.00681, 2015. `http://arxiv.org/pdf/1504.00681.pdf`.

[11] A. Kumar, R. Manokaran, M. Tulsiani, and N. K. Vishnoi. On LP-based approximability for strict CSPs. In *Proc. SODA*, pages 1560–1573, 2011.

[12] G. Kun, R. O'Donnell, S. Tamaki, Y. Yoshida, and Y. Zhou. Linear programming, width-1 CSPs, and robust satisfaction. In *Proc. ITCS*, pages 484–495, 2012.

[13] K. Makarychev and Y. Makarychev. Approximation algorithm for non-boolean Max-*k*-CSP. *Theory of Computing*, 10:341–358, 2014.

[14] E. Mossel. Gaussian bounds for noise correlation of functions. *GAFA*, 19:1713–1756, 2010.

[15] E. Mossel, R. O'Donnell, and K. Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. *Annals of Mathematics*, 171(1):295–341, 2010.

[16] P. Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proc. STOC*, pages 245–254, 2008.

[17] M. J. Serna, L. Trevisan, and F. Xhafa. The (parallel) approximability of non-boolean satisfiability problems and restricted integer programming. In *Proc. STACS*, pages 488–498, 1998.

[18] L. Trevisan. Parallel approximation algorithms by positive linear programming. *Algorithmica*, 21(1):72–88, 1998.

[19] C. Wenner. Circumventing d-to-1 for approximation resistance of satisfiable predicates strictly containing parity of width at least four. *Theory of Computing*, 9:703–757, 2013.

A LP Integrality Gap for UNIQUEGAMES

A simple probabilistic construction shows that for any $L \geq 2$ and $\delta > 0$, there is a $(1,\, (1+\delta)/L\,)$-integrality gap for the standard LP relaxation of UNIQUEGAMES on label set $[L]$.

Our randomized instance is on the n-vertex clique with uniform and normalized edge weights, where the bijective constraint for each edge is chosen uniformly and independently at random. Consider a solution to the LP relaxation in which $x_{v\ell} = 1/L$ for each vertex v and label ℓ, and $y_{e\bar{\ell}} = 1/L$ for each edge $e = (u, v)$ and $\bar{\ell} = (\ell_u, \ell_v)$ which is a satisfying assignment for the bijective constraint C_e. It is easy to see that this is a feasible solution with an LP objective of 1.

On the other hand, any fixed labeling to the vertices satisfies an edge independently with probability $1/L$, over the choice of the $\binom{n}{2}$ constraints. Thus, by Chernoff bound, the probability that a given labeling satisfies more than $(1+\delta)/L$ fraction of edges is at most,

$$p^* := \exp\left(-\frac{\delta^2 n(n-1)}{6L}\right).$$

Since the total number of possible labeling is L^n, we can choose n large enough (depending only on L and δ) so that $p^* L^n < 1$, ensuring the existence of the desired integrality gap.

ECCC ISSN 1433-8092

http://eccc.hpi-web.de