

Global Cardinality Constraints Make Approximating Some Max-2-CSPs Harder

Per Austrin 

KTH Royal Institute of Technology, Stockholm, Sweden
austrin@kth.se

Aleksa Stanković 

KTH Royal Institute of Technology, Stockholm, Sweden
aleksas@kth.se

Abstract

Assuming the Unique Games Conjecture, we show that existing approximation algorithms for some Boolean Max-2-CSPs with cardinality constraints are optimal. In particular, we prove that Max-Cut with cardinality constraints is UG-hard to approximate within ≈ 0.858 , and that Max-2-Sat with cardinality constraints is UG-hard to approximate within ≈ 0.929 . In both cases, the previous best hardness results were the same as the hardness of the corresponding unconstrained Max-2-CSP (≈ 0.878 for Max-Cut, and ≈ 0.940 for Max-2-Sat).

The hardness for Max-2-Sat applies to monotone Max-2-Sat instances, meaning that we also obtain tight inapproximability for the Max- k -Vertex-Cover problem.

2012 ACM Subject Classification Theory of computation \rightarrow Problems, reductions and completeness; Mathematics of computing \rightarrow Approximation algorithms; Theory of computation \rightarrow Approximation algorithms analysis

Keywords and phrases Constraint satisfaction problems, global cardinality constraints, semidefinite programming, inapproximability, Unique Games Conjecture, Max-Cut, Max-2-Sat

Digital Object Identifier 10.4230/LIPIcs.APPROX-RANDOM.2019.24

Category APPROX

Related Version A full version of the paper is available at <https://arxiv.org/abs/1907.04165>.

Funding Research supported by the Approximability and Proof Complexity project funded by the Knut and Alice Wallenberg Foundation.

Acknowledgements The authors thank Johan Håstad for helpful suggestions and comments on the manuscript. We also thank anonymous reviewers for their helpful remarks.

1 Introduction

Constraint satisfaction problems (CSPs) are one of the most fundamental objects studied in complexity theory. An instance of a CSP has a set of variables taking values over a certain domain and a set of constraints on tuples of these variables as an input. Probably the best known CSP is 3-Sat, in which the constraints are clauses, each clause is a disjunction of at most three literals, and each literal is either a variable or negation of a variable. In the satisfiability version of CSP problems, we are interested whether there is an assignment to the variables which satisfies all the constraints. Hardness of deciding satisfiability of CSPs is well understood, due to the dichotomy theorem [32] of Schaefer which shows that each CSP with variables taking values in a Boolean domain is either in P or NP-complete, and due to the more recent results of Bulatov [8] and Zhuk [35] which settle this question on general domains.



© Per Austrin and Aleksa Stanković;

licensed under Creative Commons License CC-BY

Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques (APPROX/RANDOM 2019).

Editors: Dimitris Achlioptas and László A. Végh; Article No. 24; pp. 24:1–24:17

Leibniz International Proceedings in Informatics



Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Another well-studied version is the Max-CSP, which is the optimization version in which we are interested in maximizing the number of constraints satisfied. This type of problem is NP-hard in most cases and we typically settle with finding a good estimate of the optimal solution, for which we rely on approximation algorithms. A common example of a constraint satisfaction problem in this setting is Max-Cut, in which the input consists of a graph G , and the goal is to partition the vertices into two sets such that the number of edges between the two parts is maximized. Approximability of Max-CSPs has been a major research topic which inspired many influential breakthroughs. One of the first surprising results was an algorithm of Goemans and Williamson [13], which uses semidefinite programming (SDP) to approximate the optimal solution to within a constant of $\alpha_{GW} \approx 0.878$. The SDP approach is also useful in approximating many other well known Max-CSPs, such as Max-3-Sat [19] within a constant of $7/8$ and Max-2-Sat [23] within $\alpha_{LLZ} \approx 0.9401$.

On the hardness of approximation side, the first NP-hardness results are based on the celebrated PCP theorem [1, 2] which provided a strong starting point for studying inapproximability. For example, a direct corollary of the PCP theorem shows that the Max-3-Sat problem cannot be approximated within $1 - \delta$ for some universal constant $\delta > 0$. By using the PCP theorem and parallel repetition [31] as a starting point, Håstad [17] proved optimal inapproximability for Max-3-Sat by showing that it cannot be approximated better than $7/8 + \epsilon$ for any $\epsilon > 0$.

However, despite further works relying on the similar techniques which improved our understanding of inapproximability for several additional CSPs, the progress on closing the gap between the best algorithm and the best hardness was at a standstill for some fundamental problems such as Max-Cut, until the Unique Games Conjecture (UGC) was introduced by Khot [20]. In particular, by assuming the UGC, optimality of the α_{GW} -approximation algorithm for Max-Cut and the α_{LLZ} -approximation algorithm for Max-2-Sat was shown in [21, 26] and [3], respectively. The strength of semidefinite programming for approximating Max-CSPs was corroborated in a breakthrough result of Raghavendra [28], which showed that assuming the UGC, a certain SDP relaxation achieves optimal approximation ratios for all Max-CSPs.

Locality of the constraints was of crucial importance in studying CSPs and Max-CSPs since their inception. Therefore, it is not a surprise that typical techniques fail when we work with CSPs for which feasible assignments need to satisfy some additional global constraints, and these problems almost always become harder. For example, while the satisfiability of a 2-Sat instance can be checked by a straightforward algorithm, Guruswami and Lee recently showed [14] that when the satisfying assignment needs to have exactly half of its variables set to true, this problem becomes NP-hard. Hardness of deciding satisfiability of CSPs in which we prescribe how many variables are assigned to certain values is well understood due to the dichotomy theorem of Bulatov and Marx [9], which shows that these problems are either NP-hard or in P, and gives a simple classification. Another type of global constraint is studied by Brakensiek et al. [7], who consider hardness of deciding CSPs in presence of modular constraints, which restrict cardinality of values in an assignment modulo a natural number M .

In this paper we are interested in optimization variants of CSPs with global cardinality constraints, i.e., constraints which specify the number of occurrences of each value from the domain in the assignment. We refer to these problems as CC-Max-CSPs. It is not hard to see that these problems are at least as hard to approximate as their unconstrained counterparts. CC-Max-CSPs have been actively studied in the past. For example the Max-Bisection problem, i.e., Max-Cut where the two partitions need to be of the same size, has been of

a particular interest, with a series of papers [12], [34],[16], [11], [30] obtaining improved approximation algorithms, until the most recent result which achieves an approximation ratio of 0.8776 [4], which is only $\approx 10^{-3}$ below the UG-hardness bound α_{GW} . The state-of-the-art algorithm [30] for the more general CC-Max-Cut problem achieves an approximation ratio of $\alpha_{cut}^{cc} \approx 0.858$. Another related CC-Max-CSP actively studied is CC-Max-2-Sat, and its monotone variant (a version in which negated literals are not allowed) Max- k -VC¹. The best algorithm [30] up to date for general CC-Max-2-Sat achieves an approximation ratio of α_{2sat}^{cc} , where $\alpha_{2sat}^{cc} \approx 0.929$, which improved on a series of increasingly stronger algorithms presented in [33], [6], and [18]. Manurangsi [25] showed that it is UG-hard to approximate Max- k -VC within a factor $\alpha_{AKS} \approx 0.944$ (note that this is slightly larger than the hardness of $\alpha_{LLZ} \approx 0.940$ for general Max-2-Sat).

Yet another well-studied CC-Max-CSP is the *Densest k -Subgraph* (Max- k -DS) problem, in which we are given a graph and the objective is to find a maximally dense induced subgraph on k vertices. Analogously to the Max- k -VC problem, Max- k -DS can be viewed as the monotone CC-Max-2-And problem. Max- k -DS is qualitatively very different from the previously discussed problems. It is not known to be approximable within a constant factor, and is in fact known to be hard to approximate to within almost polynomial factors assuming the Exponential Time Hypothesis [24], or to within any constant factor assuming the Small-Set Expansion Hypothesis [29].

Obtaining tight approximability results for CC-Max-CSPs presents an important research topic. Qualitatively, it is also interesting to determine whether adding a cardinality constraint to a non-trivial Max-CSP makes approximation strictly harder. For example, we know that CC-Max-2-Sat is as hard as Max-2-Sat, but it is still conceivable that they are equally hard. In particular, it would be interesting to answer the following question:

“Can CC-Max-2-Sat be approximated within α_{LLZ} ?”

So far the only result in this direction comes from [4] which shows that the “bisection version” (where the cardinality constraint is that exactly half of the variables must be set to true) of CC-Max-2-Sat can be approximated within α_{LLZ} . However, the approach taken in that algorithm does not immediately extend to handle general cardinality constraints. A similar question arises for the CC-Max-Cut problem, but here even the basic Max-Bisection problem is not known to be approximable within the Max-Cut constant $\alpha_{GW} \approx 0.878$. As far as we are aware, prior to this paper, the only examples of cardinality-constrained Max-CSPs being harder than their unconstrained counterparts were examples where the unconstrained version is easy (e.g. unconstrained Max- k -VC is monotone Max-2-Sat, and unconstrained Max- k -DS is monotone Max-2-And, which are both trivial).


Our Results

In this paper, we answer the above question negatively, by giving improved UG-hardness results for CC-Max-Cut and Max- k -VC.

► **Theorem 1.** *For every $\varepsilon > 0$, CC-Max-Cut is UG-hard to approximate within $\beta_{cut}^{cc} + \varepsilon$, where $\beta_{cut}^{cc} \approx 0.858$.*

► **Theorem 2.** *For every $\varepsilon > 0$, Max- k -VC is UG-hard to approximate within $\beta_{vc}^{cc} + \varepsilon$, where $\beta_{vc}^{cc} \approx 0.929$.*

¹ Max- k -VC is an abbreviation for maximum k vertex cover, in which we are given a graph and the task is to select a subset of k vertices covering as many of the edges as possible.



■ **Figure 1** Hardness ratio (red) vs. approximation constant (blue) as a function of cardinality constraint q for CC-Max-Cut (left), as well as Max- k -VC (right). We also use \times to highlight the approximation ratio known for Max-Bisection on the left plot, while \times represents the optimal approximation constant for the CC-Max-2-Sat problem in the special case $q = 1/2$.

Note that since CC-Max-Cut and Max- k -VC are special cases of CC-Max-2-Lin and CC-Max-2-Sat respectively, the corresponding hardness results apply to the latter problems as well.

The constants β_{vc}^{cc} and β_{cut}^{cc} are calculated numerically and their estimated values match the constants α_{2sat}^{cc} and α_{cut}^{cc} , which are the approximation ratios for corresponding problems achieved by the algorithm of Raghavendra and Tan [30]. We provide even stronger evidence that these constants match each other, by showing that β_{vc}^{cc} and β_{cut}^{cc} are calculated as minima of the same functions used for calculating their counterparts α_{vc}^{cc} and α_{cut}^{cc} , but over a slightly more restricted domain.

Moreover, in Section 3 we give refined statements of Theorem 1 and Theorem 2 which describe inapproximability of these problems as a function of the cardinality constraint $q \in (0, 1)$, which specifies the fraction of variables that need to be set to true. For now, we provide a visualization of these results in Figure 1.

Overview of proof ideas

The main observation behind the hardness results is that the reduction used to prove hardness of approximation for the Independent Set and Vertex Cover problems in bounded degree graphs [5] gives very strong soundness guarantees. In particular it shows that in the “no” case of the reduction, all induced subgraphs of the graph contain many edges, which in turn gives useful upper bounds on the number of edges cut by a bipartition of a given size, or the number of edges covered by a subgraph. This is also how [25] obtained the previous hardness of ≈ 0.944 for Max- k -VC. Thus our results use essentially the same reduction as [5] (which is in turn similar to the reduction for Max-Cut [21]). Note however that even though the graph produced by that reduction has a small vertex cover in the “yes” case, using that small vertex cover is not necessarily the best solution for the Max- k -VC problem on the graph. In particular for $q < 1/2$, it makes more sense to instead use the large independent set as the Max- k -VC solution in the yes case (the intuition being that since it is independent, it covers many edges relative to its size).

Another difference is that we have somewhat greater flexibility in choosing the noise distribution of our “dictatorship test” (the key component of essentially all UG-hardness results) The reason is that for Independent Set/Vertex Cover, the reduction needs “perfect completeness”, i.e., in the “yes” case it needs to produce graphs with large independent

sets/small vertex covers, whereas for e.g. Max- k -VC we are perfectly happy with graphs where there are sets of size k covering many, but not necessarily all, edges. This increased flexibility turns out to improve the hardness ratios for some range of the cardinality constraint q . For example, for the CC-Max-Cut problem with $q = 1/2$, this allows us to recover the α_{GW} -hardness for the Max-Bisection problem using the same reduction. However, at q further away from $1/2$, and in particular at the local minima in Figure 1, it turns out that this flexibility does not help. Thus in the global minimum at $q \approx 0.365$ for Max- k -VC, the reduction outputs a graph with a large independent set containing a q fraction of the vertices, and choosing that independent set is the optimal solution for the Max- k -VC instance. Similarly, at the local minimum with $q > 1/2$, the optimal solution to the Max- k -VC instance in the yes case is to pick an actual vertex cover of size q , and this point of the curve corresponds exactly to the hardness of 0.944 from [25].

Organization

This paper is organized as follows. In Section 2 we fix the notation, recall some well-known facts, and formally introduce the problems of interest. In Section 3 we give our improved inapproximability results. In Section 4 we give a brief overview of the algorithm of Raghavendra and Tan [30] in order to observe that the hardness ratios we get match the approximation ratios of the algorithm. Finally, in Section 5 we propose some possible directions for future research.

2 Preliminaries

2.1 Notational Conventions

In this paper we work with undirected (multi)graphs $G = (V, E)$. For a set $S \subseteq V$ of vertices we use S^c to denote its complement $S^c = V \setminus S$, and write $U \sqcup V$ for a disjoint union of sets U and V . The graphs are both edge and vertex weighted and the weights of vertices and edges are given by functions $w: V \rightarrow [0, 1]$, and $w: E \rightarrow [0, 1]$. For subsets $S \subseteq V$ and $K \subseteq E$ we interpret $w(S)$ and $w(K)$ as the sum of weights of vertices contained in S and edges in K , respectively. Furthermore, weights are normalized so that $w(V) = w(E) = 1$ and the weight of each vertex equals half the weight of all edges adjacent to it. Therefore, the weights of edges and vertices can be interpreted as probability distributions, and sampling a vertex with probability equal to its weight is the same as sampling an edge and then sampling one of its endpoints with probability $1/2$. For $S, T \subseteq V$, we write $w(S, T)$ for the total weight of edges from E which have one endpoint in S , and other in T . Note that, since we work with undirected graphs, the order of endpoints is not important, and therefore $w(S, T) = w(T, S)$. In other words, the weight of an edge $e = (u, v)$ contributes to $w(S, T)$ if either $(u, v) \in T \times S$ or $(u, v) \in S \times T$. We also have the identity

$$w(S, V) = w(S) + \frac{1}{2}w(S, S^c). \quad (1)$$

The set of all neighbours of a vertex v including v is denoted by $N(v)$, and the set of all neighbours of a set $S \subseteq V$ including S is denoted by $N(S)$. Let us also introduce the following definition.

► **Definition 3.** *A graph G is (q, ε) -dense if every subset $S \subseteq V$ with $w(S) = q$ satisfies $w(S, S) \geq \varepsilon$.*

We use $\phi(x) = \frac{1}{\sqrt{2\pi}}e^{-x^2/2}$ to denote the density function of a standard normal random variable, and $\Phi(x) = \int_{-\infty}^x \phi(y)dy$ to denote its cumulative distribution function (CDF). We also work with bivariate normal random variables, and to that end introduce the following function.

► **Definition 4.** Let $\rho \in [-1, 1]$, and consider two jointly normal random variables X, Y with mean 0 and covariance matrix $\text{Cov}(X, Y) = \begin{bmatrix} 1 & \rho \\ \rho & 1 \end{bmatrix}$. We define $\Gamma_\rho: [0, 1]^2 \rightarrow [0, 1]$ as

$$\Gamma_\rho(x, y) = \Pr [X \leq \Phi^{-1}(x) \wedge Y \leq \Phi^{-1}(y)].$$

We also write $\Gamma_\rho(x) = \Gamma_\rho(x, x)$. We have the following basic lemma (for a proof see Appendix A of [4]).

► **Lemma 5.** For every $\rho \in [-1, 1]$, and every $x, y \in [0, 1]$, we have

$$\Gamma_\rho(x, y) = \Gamma_\rho(1-x, 1-y) - 1 + x + y.$$

2.2 Problem Definitions

This paper is concerned with Max-Cut, Max-2-Lin, Max-2-Sat, and Max- k -VC problems with cardinality constraints. Let us give the definitions of these problems as integer optimization programs now. In these definitions instead of $\{0, 1\}$ we represent Boolean domain as $\{-1, 1\}$, and for that reason instead of cardinality constraint q we consider a *balance constraint* $r = 1 - 2q$.

► **Definition 6.** An instance \mathcal{F} of the cardinality constrained Max-2-Lin (CC-Max-2-Lin) problem with balance constraint $r \in (-1, 1)$ over variables $X = \{x_1, \dots, x_n\}$ taking values in $\{-1, 1\}$ is given by the following integer optimization program

$$\begin{aligned} \max \quad & \sum_{(i,j)=e_\ell \in E} \frac{1 + P_\ell x_i x_j}{2}, \\ \text{s.t.} \quad & \sum_{i \in V} x_i = nr, \end{aligned}$$

where $P_\ell \in \{-1, 1\}$ and the term $(1 + P_\ell x_i x_j)/2$ corresponds to the XOR constraint $x_i x_j = P_\ell$. In case $P_\ell = -1$ for all ℓ , the integer optimization program is an instance of CC-Max-Cut problem.

► **Definition 7.** An instance \mathcal{F} of the cardinality constrained Max-2-Sat (CC-Max-2-Sat) problem with balance constraint $r \in (-1, 1)$ over variables $X = \{x_1, \dots, x_n\}$ taking values in $\{-1, 1\}$ is given by the following integer optimization program

$$\begin{aligned} \max \quad & \sum_{(i,j)=e_\ell \in E} \frac{3 + P_\ell^1 x_i + P_\ell^2 x_j + P_\ell^3 x_i x_j}{4}, \\ \text{s.t.} \quad & \sum_{i \in V} x_i = nr, \end{aligned}$$

where $(P_\ell^1, P_\ell^2, P_\ell^3) \in \{(-1, -1, -1), (1, -1, 1), (-1, 1, 1), (1, 1, -1)\}$ corresponds to one of the four possible clauses

$$x_i \vee x_j, \quad \neg x_i \vee x_j, \quad x_i \vee \neg x_j, \quad \neg x_i \vee \neg x_j.$$

In case $(P_\ell^1, P_\ell^2, P_\ell^3) = (-1, -1, -1)$ for all ℓ , the integer optimization program is an instance of Max- k -VC problem.

The objective in the problems given by Definitions 6 and 7 is to find an assignment $z: X \rightarrow \{-1, 1\}$ which satisfies a (hard) global cardinality constraint and maximizes the number of satisfied soft constraints represented by the objective function. For an assignment z that satisfies global constraint of an instance \mathcal{F} we use $\text{Val}_z(\mathcal{F})$ to denote the value of the objective function under the assignment z . Furthermore, we use

$$\text{OptVal}(\mathcal{F}) = \max_{\substack{z: X \rightarrow \{-1, 1\} \\ \sum_{x \in X} z(x) = rn}} \text{Val}_z(\mathcal{F})$$

to denote the maximum value of the objective function over all assignments z satisfying the cardinality constraint.

The starting point of the hardness results in this paper is the Unique Games problems, which is defined as follows.

► **Definition 8.** A Unique Games instance $\Lambda = (\mathcal{U}, \mathcal{V}, \mathcal{E}, \Pi, [L])$ consists of an unweighted bipartite multigraph $(\mathcal{U} \sqcup \mathcal{V}, \mathcal{E})$, a set $\Pi = \{\pi_e: [L] \rightarrow [L] \mid e \in \mathcal{E} \text{ and } \pi_e \text{ is a bijection}\}$ of permutation constraints, and a set $[L]$ of labels. The value of Λ under the assignment $z: \mathcal{U} \sqcup \mathcal{V} \rightarrow [L]$ is the fraction of edges satisfied, where an edge $e = (u, v), u \in \mathcal{U}, v \in \mathcal{V}$ is satisfied if $\pi_e(z(u)) = z(v)$. We write $\text{Val}_c(\Lambda)$ for the value of Λ under z , and $\text{Opt}(\Lambda)$ for the maximum possible value over all assignments z .

The Unique Games Conjecture [20] can be formulated as follows ([22], Lemma 3.4).

► **Conjecture 9** (Unique Games Conjecture). For every constant $\gamma > 0$ there is a sufficiently large $L \in \mathbb{N}$, such that for a Unique Games instance $\Lambda = (\mathcal{U}, \mathcal{V}, \mathcal{E}, \Pi, [L])$ with a regular bipartite graph $(\mathcal{U} \sqcup \mathcal{V}, \mathcal{E})$, it is NP-hard to distinguish between

- $\text{Opt}(\Lambda) \geq 1 - \gamma$,
- $\text{Opt}(\Lambda) \leq \gamma$.

2.3 Analysis of Boolean Functions

One of the ubiquitous tools in the hardness of approximation area is Fourier analysis of Boolean functions. We now recall some of the well-known facts which are used in the paper. For a more detailed study, we refer to [27].

For $q \in [0, 1]$ and $n \in \mathbb{N}$ we write $\pi_q: \{0, 1\} \rightarrow [0, 1]$ for the probability distribution given by $\pi_q(1) = q, \pi_q(0) = 1 - q$. We also write $\pi_q^{\otimes n}$ for the probability distribution on n -bit strings $x \in \{0, 1\}^n$ where each bit is distributed according to π_q , independently. We use $L^2(\pi_q^{\otimes n})$ to denote the space of random variables $f: \{0, 1\}^n \rightarrow \mathbb{R}$ over the probability space $(\{0, 1\}^n, P(\{0, 1\}^n), \pi_q^{\otimes n})$, and interpret $\mathbf{E}[f]$ and $\mathbf{Var}[f]$ as expectation and variance of $f(X)$ when the X is drawn from $\pi_q^{\otimes n}$. Depending on context, the elements of $L^2(\pi_q^{\otimes n})$ will be interpreted as functions as well.

Let us now introduce some of the common objects used in the study of Boolean functions.

► **Definition 10.** Consider a function $f \in L^2(\pi_q^{\otimes n})$ and $i \in \{1, \dots, n\}$. The influence $\mathbf{Inf}_i[f]$ of the i -th argument on f is defined as

$$\mathbf{Inf}_i[f] = \mathbf{E}_{x \sim \pi_q^{\otimes n}} [\mathbf{Var}_{\tilde{x}_i \sim \pi_q} [f(x_1, \dots, x_{i-1}, \tilde{x}_i, x_{i+1}, \dots, x_n)]]$$

Minimal correlation between two q -biased bits is $\max(-q/(1 - q), -(1 - q)/q)$. For notational convenience, let us introduce the function κ which assigns to each value $q \in (0, 1)$ an interval $I \subseteq (-1, 0)$ as

$$\kappa(q) = \begin{cases} [-q/(1 - q), 0), & \text{if } q < 1/2, \\ (-1, 0), & \text{if } q = 1/2, \\ [-(1 - q)/q, 0), & \text{if } q > 1/2. \end{cases}$$

► **Definition 11.** For a fixed $x \in \{0, 1\}$, $q \in (0, 1)$ and $\rho \in \kappa(q)$ we write $y \sim N_\rho(x)$ to indicate that y is a ρ -correlated copy of x . In particular each bit y_i is equal to 1 with probability $q + \rho(1 - q)$ if $x_i = 1$, and $y_i = 1$ with probability $q - \rho q$ when $x_i = 0$, independently.

► **Definition 12.** Consider $q \in (0, 1)$ and $\rho \in \kappa(q)$. The noise operator $T_\rho: L^2(\pi_q^{\otimes n}) \rightarrow L^2(\pi_q^{\otimes n})$ is defined as

$$T_\rho f(x) = \mathbf{E}_{y \sim N_\rho(x)}[f(y)].$$

The following lemma gives a useful bound on the number of influential variables of $T_\rho f$.

► **Lemma 13.** Consider $q \in (0, 1)$, a function $f \in L^2(\pi_q^{\otimes n})$, and $\rho \in \kappa(q)$. Then, for any $\tau > 0$ we have that

$$|\{i \in [n] \mid \mathbf{Inf}_i[T_\rho f] \geq \tau\}| \leq \frac{\mathbf{Var}[f]}{\tau e \ln(1/|\rho|)}.$$

For a proof we refer to Lemma 3.4 of [15]. We also need to introduce the notion of noise stability, defined as follows.

► **Definition 14.** Let $q \in (0, 1)$, $\rho \in \kappa(q)$ and $f \in L^2(\pi_q^{\otimes n})$. The noise stability of function f at ρ is defined as

$$\mathbb{S}_\rho = \mathbf{E}[f \cdot T_\rho f].$$

Let us also recall the following variant of the “Majority is Stablest” theorem in the form that appeared in [5], and which follows from Theorem 3.1 in [10].

► **Theorem 15.** Let $q \in (0, 1)$ and $\rho \in \kappa(q)$. Then for any $\varepsilon > 0$, there exist $\tau > 0$ and $\delta > 0$ such that for every function $f \in L^2(\pi_q^{\otimes n})$, $f: \{-1, 1\}^n \rightarrow [0, 1]$ that satisfies

$$\max_{i \in [n]} \mathbf{Inf}_i[T_{1-\delta} f] \leq \tau,$$

we have

$$\mathbb{S}_\rho(f) \geq \Gamma_\rho(\mathbf{E}[f]) - \varepsilon.$$

3 Hardness Reduction

In this section we give our main hardness reduction. As discussed in the introduction, it is a generalization of the reduction of Theorem III.1 from [5].

► **Theorem 16.** For every $q \in (0, 1)$, $\varepsilon > 0$, and $\rho \in \kappa(q)$, there exists a $\gamma > 0$ and a reduction from Unique Games instances $\Lambda = (\mathcal{U}, \mathcal{V}, \mathcal{E}, \Pi, [L])$ to weighted multigraphs $G = (V, E)$ with the following properties:

- Completeness: If $\text{Opt}(\Lambda) \geq 1 - \gamma$, then there is a set $S \subseteq V$ such that $w(S) = q$ and $w(S, S^c) \geq 2q(1 - q)(1 - \rho) - 2\gamma$.
- Soundness: If $\text{Opt}(\Lambda) \leq \gamma$, then for every $r \in [0, 1]$, G is $(r, \Gamma_\rho(r) - \varepsilon)$ -dense.

Moreover, the running time of the reduction is polynomial in $|\mathcal{U}|, |\mathcal{V}|, |\mathcal{E}|$, and exponential in L .

Proof. Let $\nu: \{0, 1\}^2 \rightarrow [0, 1]$ be the probability distribution over two ρ -correlated q -biased bits. In other words, letting $t = (q - q^2)(1 - \rho)$, we have

$$\nu(0, 0) = 1 - q - t, \quad \nu(0, 1) = \nu(1, 0) = t, \quad \nu(1, 1) = q - t.$$

Let us now describe how the multigraph G can be constructed from Λ . We define the vertex set of G to be $V = \mathcal{V} \times \{0, 1\}^L = \{(v, x) \mid v \in \mathcal{V}, x \in \{0, 1\}^L\}$. In particular, for every vertex $v \in \mathcal{V}$ we create 2^L vertices of G , which we identify with L -bit strings in $\{0, 1\}^L$. We also write v^x for a vertex (v, x) of the graph G . The weights of vertices in G are given by

$$w(v^x) = \frac{1}{|\mathcal{V}|} \pi_q^{\otimes L}(x). \quad (2)$$

The edges of G are constructed in the following way. For every $u \in \mathcal{U}$, and for every two $v_1, v_2 \in N(u)$, we create an edge between vertices v_1^x, v_2^y with weight

$$\frac{1}{|\mathcal{U}|D^2} \nu^{\otimes L}(x \circ \pi_{e_1}, y \circ \pi_{e_2}), \quad \text{where } e_1 = (u, v_1), \quad e_2 = (u, v_2).$$

Expressed formally, the edge set E is

$$E = \{(e_1^x, e_2^y) \mid e_1 = (u, v_1), e_2 = (u, v_2), u \in \mathcal{U}, v_1, v_2 \in \mathcal{V}, x, y \in \{0, 1\}^L\}.$$

Since the marginal of the distribution ν over either the first or the second argument is a q -biased distribution on $\{0, 1\}^L$, the weight of all edges adjacent to a vertex v^x equals two times the weight of the vertex v^x . Furthermore, it is trivial to check that $w(V) = w(E) = 1$. The number of vertices in G is $|\mathcal{V}|2^L$, and the number of edges is $|\mathcal{U}|D^22^L$, so the construction is indeed polynomial in $|\mathcal{U}|, |\mathcal{V}|$ and $|\mathcal{E}|$.

Let us now prove completeness and soundness.

Completeness: Since $\text{Opt}(\Lambda) \geq 1 - \gamma$, there is a labeling $z: \mathcal{U} \sqcup \mathcal{V} \rightarrow [L]$ such that $\text{Val}_z(\Lambda) \geq 1 - \gamma$. Consider a set S given by

$$S = \{v^x \in V \mid x_{z(v)} = 1\}.$$

The weight of the set S is obviously q . Let us consider a set consisting of pairs of edges in \mathcal{E} which have a common vertex in \mathcal{U} , i.e. the set

$$\hat{E} = \{(e_1, e_2) \in \mathcal{E} \times \mathcal{E} \mid e_1 = (u, v_1), e_2 = (u, v_2), u \in \mathcal{U}, v_1, v_2 \in \mathcal{V}\},$$

and its subset \hat{E}_{good} consisting of edge pairs which are satisfied under the assignment z , or formally

$$\hat{E}_{\text{good}} = \{(e_1, e_2) \in \hat{E} \mid e_1 = (u, v_1), e_2 = (u, v_2), z(u) = \pi_{e_1}^{-1}(z(v_1)) = \pi_{e_2}^{-1}(z(v_2))\},$$

Since at least fraction $1 - \gamma$ of edges in \mathcal{E} are satisfied under z , at least fraction $(1 - \gamma)^2$ of edge pairs in \hat{E} is satisfied under z , i.e. $|\hat{E}_{\text{good}}| \geq (1 - \gamma)^2 |\hat{E}|$. For every $(e_1, e_2) \in \hat{E}_{\text{good}}$, $e_1 = (u, v_1), e_2 = (u, v_2)$, the edges between S and S^c created through the pair of edges (e_1, e_2) have the total weight of

$$\begin{aligned} \frac{1}{|\mathcal{U}|D^2} \Pr_{(x,y) \sim \nu^{\otimes L}} \left[(x \circ \pi_{e_1}^{-1})_{z(v_1)} \neq (y \circ \pi_{e_2}^{-1})_{z(v_2)} \right] &= \frac{1}{|\mathcal{U}|D^2} \Pr_{(x,y) \sim \nu^{\otimes L}} [x_{z(u)} \neq y_{z(u)}] \\ &= \frac{1}{|\mathcal{U}|D^2} (\nu(0, 1) + \nu(1, 0)) = \frac{1}{|\mathcal{U}|D^2} 2t. \end{aligned}$$

Therefore, we have $w(S, S^c) \geq 2t(1 - \gamma)^2 \geq 2q(1 - q)(1 - \rho) - 2\gamma$.

Soundness: Let us assume by contradiction that G is not $(r, \Gamma_\rho(r) - \varepsilon)$ -dense, and therefore that there is a set $S \subseteq V$ of weight $w(S) = r$ for which $w(S, S) < \Gamma_\rho(r) - \varepsilon$. For each $v \in \mathcal{V}$, let us define a function $S_v \in L^2(\pi_q^{\otimes L})$ to be the indicator function of S restricted

24:10 Global Cardinality Constraints Make Approximating Some Max-2-CSPs Harder

to the vertex v . In particular, we have that $S_v(x) = 1$ if and only if $v^x \in S$. Furthermore, for all $u \in \mathcal{U}$ let us define $S_u \in L^2(\pi_q^{\otimes L})$ as

$$S_u(x) = \mathbf{E}_{\substack{e=(u,v), \\ v \in N(u)}} [S_v(x \circ \pi_e^{-1})].$$

We have that

$$\begin{aligned} w(S, S) &= \mathbf{E}_{\substack{u \in \mathcal{U}, \\ e_1=(u,v_1), e_2=(u,v_2) \\ v_1, v_2 \in N(u)}} \left[\mathbf{E}_{(x,y) \sim \nu^{\otimes L}} [S_{v_1}(x \circ \pi_{e_1}^{-1}) S_{v_2}(y \circ \pi_{e_2}^{-1})] \right] \\ &= \mathbf{E}_{\substack{u \in \mathcal{U}, \\ (x,y) \sim \nu^{\otimes L}}} \left[\mathbf{E}_{\substack{e_1=(u,v_1), e_2=(u,v_2) \\ v_1, v_2 \in N(u)}} [S_{v_1}(x \circ \pi_{e_1}^{-1}) S_{v_2}(y \circ \pi_{e_2}^{-1})] \right] \\ &= \mathbf{E}_{u \in \mathcal{U}} \left[\mathbf{E}_{(x,y) \sim \nu^{\otimes L}} [S_u(x) S_u(y)] \right] = \mathbf{E}_{u \in \mathcal{U}} \left[\mathbf{E}_{x \sim \pi_q^{\otimes L}} [S_u(x) T_\rho S_u(x)] \right] = \mathbf{E}_{u \in \mathcal{U}} [\mathbb{S}_\rho(S_u)]. \end{aligned}$$

Let us define $\mu_u = \mathbf{E}_{x \sim \pi_q^{\otimes L}} [S_u(x)]$, and remark that due to regularity of Λ we have $\mathbf{E}_{u \in \mathcal{U}} [S_u] = r$. We claim that there is a set $\mathcal{U}' \subseteq \mathcal{U}$, $|\mathcal{U}'| \geq \varepsilon |\mathcal{U}|/2$ such that for every $u \in \mathcal{U}'$ we have $\mathbb{S}_\rho(S_u) < \Gamma_\rho(\mu_u) - \varepsilon/2$. Otherwise, we reach a contradiction by noticing that

$$\begin{aligned} \Gamma_\rho(r) - \varepsilon > w(S, S) &= \mathbf{E}_{u \in \mathcal{U}} [\mathbb{S}_\rho(S_u)] \geq (1 - \varepsilon/2) \left(\mathbf{E}_{u \in \mathcal{U}} [\Gamma_\rho(\mu_u)] - \varepsilon/2 \right) \\ &\geq \mathbf{E}_{u \in \mathcal{U}} [\Gamma_\rho(\mu_u)] - \varepsilon \geq \Gamma_\rho(r) - \varepsilon, \end{aligned}$$

where in the last inequality we used the fact that Γ_ρ is convex.

By Theorem 15 there is $\tau > 0$ and $\delta > 0$ such that for every $u \in \mathcal{U}'$ there is a significant coordinate $i \in [L]$ for which $\mathbf{Inf}_i[T_{1-\delta} S_u] \geq \tau$. For each $u \in \mathcal{U}'$ and for its significant coordinate i , by using the fact that \mathbf{Inf}_i is convex and Markov's inequality we conclude that for at least $\tau/2$ of $v \in N(u)$ we have

$$\mathbf{Inf}_{\pi_e(i)} [T_{1-\delta} S_v] \geq \tau/2, \quad e = (u, v).$$

For each $v \in \mathcal{V}$ let $[L]_v \subseteq [L]$ denote a set of labels defined by

$$[L]_v = \{i \in [L] \mid \mathbf{Inf}_i [T_{1-\delta} S_v] \geq \tau/2\}.$$

By Lemma 13 we have that $|[L]_v| \leq \frac{2}{\tau \varepsilon \ln(1/(1-\delta))}$. Let us now pick an assignment $z: \mathcal{U} \sqcup \mathcal{V} \rightarrow [L]$ of Λ using the following randomized procedure. For each $v \in \mathcal{V}$, pick $i \in [L]_v$ randomly, and set $z(v) = i$. If $[L]_v = \emptyset$, we pick $i \in [L]$ randomly. Then, for each $u \in \mathcal{U}$, we set $z(u) = i$ for the i that maximizes the number of edges satisfied. From the previous discussion we conclude that this labeling satisfies $\Omega(\varepsilon \tau^4 \ln^2(1/(1-\delta)))$ of constraints of Λ in expectation. But since this constant does not depend on γ this would be a contradiction if we started with a sufficiently small γ . \blacktriangleleft

3.1 Hardness for CC-Max-Cut

Now that we have proven Theorem 16, it is straightforward to prove the following theorem which gives a hardness result of CC-Max-Cut.

► **Theorem 17.** *For any $q \in (0, 1)$ and $\rho \in \kappa(q)$ it is UG-hard to approximate CC-Max-Cut with cardinality constraint q within $\beta_{cut}^{cc}(q, \rho) + \varepsilon$ where $\varepsilon > 0$ is arbitrary small and $\beta_{cut}^{cc}(q, \rho)$ is given by*

$$\beta_{cut}^{cc}(q, \rho) = \frac{1 - \Gamma_\rho(q) - \Gamma_\rho(1 - q)}{2(q - q^2)(1 - \rho)}.$$

Proof. By Theorem 16 there exists a family of multigraphs $G = (V, E)$ for which it is UG-hard to decide between the following two statements:

- There is a set $S \subseteq V, w(S) = q$, such that $w(S, S^c) \geq 2q(1 - q)(1 - \rho) - 2\gamma$.
- For any $r \in [0, 1]$ and every set $T \subseteq V, w(T) = r$ we have $w(T, T) \geq \Gamma_\rho(r) - \varepsilon$.

The second statement implies that for any $S \subseteq V, w(S) = q$, we have $w(S, S^c) = w(V, V) - w(S, S) - w(S^c, S^c) \leq 1 - \Gamma_\rho(q) - \Gamma_\rho(1 - q) + 2\varepsilon$. Therefore, by setting γ sufficiently small this shows UG-hardness of approximating CC-Max-Cut with cardinality constraint q within

$$\frac{1 - \Gamma_\rho(1 - q) - \Gamma_\rho(q)}{2q(1 - q)(1 - \rho)} + 2\varepsilon,$$

where $\varepsilon > 0$ is arbitrarily small. This reduction yields a weighted graph, which can be easily converted into an unweighted multigraph, using e.g. a simple reduction from Step 1 of Theorem 4.1. in [5]. ◀

3.2 Hardness for Max- k -VC

Next we give the hardness result for Max- k -VC.

► **Theorem 18.** *Consider $q \in (0, 1)$ and let $\rho \in \kappa(q)$. Then, it is UG-hard to approximate Max- k -VC with cardinality constraint q within $\beta_{vc}^{cc}(q, \rho) + \varepsilon$ where $\varepsilon > 0$ is arbitrary small and $\beta_{vc}^{cc}(q, \rho)$ is given by*

$$\beta_{vc}^{cc}(q, \rho) = \frac{1 - \Gamma_\rho(1 - q)}{q(1 + (1 - q)(1 - \rho))}.$$

Proof. As we have shown in Theorem 16, there is a family of multigraphs $G = (V, E)$ for which it is UG-hard to decide between the following two statements:

- There is a set $S \subseteq V, w(S) = q$, such that $w(S, S^c) \geq 2q(1 - q)(1 - \rho) - 2\gamma$.
- For any $r \in [0, 1]$ and every set $T \subseteq V, w(T) = r$ we have $w(T, T) \geq \Gamma_\rho(r) - \varepsilon$.

By (1), the first item implies that $w(S, V) = q(1 + q(1 - q)(1 - \rho)) - \gamma$. The second statement implies that for any $S \subseteq V, w(S) = q$, we have $w(S, V) = w(V, V) - w(S^c, S^c) \leq 1 - \Gamma_\rho(1 - q) + \varepsilon$. Therefore, by letting $\gamma \rightarrow 0$ this shows UG-hardness of approximating Max- k -VC with cardinality constraint q within

$$\frac{1 - \Gamma_\rho(1 - q)}{q(1 + (1 - q)(1 - \rho))} + \varepsilon,$$

where $\varepsilon > 0$ is arbitrarily small. As in the CC-Max-Cut case, this reduction yields a weighted graph, which can be converted into an unweighted multigraph by using the reduction from [5]. ◀

3.3 Hardness as a Function of the Cardinality Constraint

As we have concluded in Theorems 17 and 18, it is UG-hard to approximate CC-Max-Cut and Max- k -VC with cardinality constraint $q \in (0, 1)$ to within

$$\beta_{cut}^{cc}(q) = \inf_{\rho \in \kappa(q)} \beta_{cut}^{cc}(q, \rho), \quad \beta_{vc}^{cc}(q) = \inf_{\rho \in \kappa(q)} \beta_{vc}^{cc}(q, \rho),$$

respectively. For a fixed q it is not clear for which ρ the functions $\beta_{cut}^{cc}(q, \cdot)$ and $\beta_{vc}^{cc}(q, \cdot)$ are minimized. For the plots of the inapproximability curves in Figure 1, the optimization over ρ was done numerically. Interestingly, numerical calculations show that the worst-case value of the cardinality constraint $q < 1/2$ (the value of q at which the hardness ratio meets the approximation ratio) is the same for Max- k -VC and CC-Max-Cut, and in particular its value is $q^* \approx 0.365$. The value of the correlation parameter ρ for which this worst-case hardness is achieved is extremal, i.e., $\rho = -q^*/(1 - q^*) \approx -0.575$. However, the local minima at $q > 1/2$ in the two curves do not occur at the same value of q . For CC-Max-Cut the curve is symmetric around $1/2$ and the minimum occurs at $1 - q^* \approx 0.635$, but for the less symmetric Max- k -VC problem it occurs at ≈ 0.574 .

Furthermore, for all $q \leq q^*$ and also for $q > 1/2$ greater than the respective local minimum, the ρ minimizing both $\beta_{cut}^{cc}(q, \rho)$ and $\beta_{vc}^{cc}(q, \rho)$ is the minimum value of $\kappa(q)$. On the other hand, when q is close to $1/2$, the best choice of ρ does not equal $\min \kappa(q)$. For example, when $q = 1/2$, the hardness we obtain for CC-Max-Cut is the same as for the Max-Cut problem, attained using the value $\rho \approx -0.689$.

4 Approximation Algorithm

In this section we recall the algorithm of Raghavendra and Tan [30], somewhat reformulated in order to obtain explicit expressions for the approximation ratios that match the hardness results we obtain. We keep the exposition at a high level and skip over certain technical details, and refer the reader interested in the details to [30] or the follow-up work [4].

In order to find a good approximation for NP-hard integer optimization problems given in Definitions 6 and 7 we use semidefinite programming (SDP) relaxations. In particular, we extend the domain of variables $\{x_i\}_{i=1}^n$ from $\{0, 1\}$ to vectors on an n -sphere, which we denote by $v_i \in S^n$. We also introduce a vector $v_0 \in S^n$ which represents the value false (corresponding value is 1 in the integer program). Then, we replace x_i by the scalar product $\langle v_0, v_i \rangle$ and $x_i x_j$ with $\langle v_i, v_j \rangle$. For example, the semidefinite relaxation of the CC-Max-Cut program is given as

$$\begin{aligned} \max \quad & \sum_{(i,j)=e \in E} \frac{1 - \langle v_i, v_j \rangle}{2}, \\ \text{s.t.} \quad & \sum_{i \in V} \langle v_i, v_0 \rangle = rn. \end{aligned}$$

Furthermore, since $|x_i - x_j| \leq |x_i - x_k| + |x_k - x_j|$, we also demand from the vectors v_i to satisfy the triangle inequalities $\|v_i - v_j\|_2^2 \leq \|v_i - v_0\|_2^2 + \|v_0 - v_j\|_2^2$. In order to relax the notation we define $\mu_i = \langle v_0, v_i \rangle$, $\rho_{ij} = \langle v_i, v_j \rangle$, and write triangle inequalities as

$$\begin{aligned} \mu_i + \mu_j + \rho_{ij} &\geq -1, & \mu_i - \mu_j - \rho_{ij} &\geq -1, \\ -\mu_i + \mu_j - \rho_{ij} &\geq -1, & -\mu_i - \mu_j + \rho_{ij} &\geq -1. \end{aligned}$$

The triples (μ_1, μ_2, ρ) satisfying triangle inequalities will be called *configurations*. We denote the set of all configurations as $\mathbf{Conf} \subseteq [-1, 1]^3$. We can solve a semidefinite program up to desired accuracy in polynomial time. Then, the main challenge is finding a *rounding algorithm* which translates the vectors $\{v_i\}_{i=0}^n$ back to $\{-1, 1\}$ so that they satisfy the balance constraint, and such that the rounding does not incur a big loss in the objective value. Raghavendra and Tan used a randomized rounding procedure, which rounds vectors

$\{v_i\}_{i=0}^n$ to ± 1 integers $\{y_i\}_{i=1}^n$ in the following way. First, let us define $w_i = v_i - \mu_i v_0$, and let² $\bar{w}_i = w_i / \|w_i\|$. Then, we draw a vector g from the Gaussian distribution $\mathcal{N}(0, I^{n+1})$ and set the values of \bar{y}_i as

$$\bar{y}_i = \begin{cases} 1 & \text{if } \langle g, \bar{w}_i \rangle \geq \Phi^{-1}\left(\frac{1-\mu_i}{2}\right), \\ -1 & \text{otherwise.} \end{cases}$$

It is trivial to check that $\mathbf{E}[\bar{y}_i] = \mu_i$, so we have $\mathbf{E}\left[\sum_{i=1}^n \bar{y}_i\right] = rn$, and therefore the solution $\{\bar{y}_i\}_{i=1}^n$ satisfies the balance constraint in expectation. Furthermore, as shown in [30], using additional levels of the Lasserre hierarchy we can guarantee that with probability $1 - \delta$ the sampled solution $\{\bar{y}_i\}_{i=1}^n$ is $O(\delta)$ -far away from satisfying the balance constraint, where $\delta > 0$ can be chosen arbitrarily small. Therefore, we can change the values of at most $O(\delta)n$ variables \bar{y}_i to get a solution y_i exactly satisfying the balance constraint, while losing only an additional small factor $O(\delta)$ in the objective value. Thus, it is sufficient to show that the objective value of the \bar{y}_i 's is large.

Consider now the SDP relaxation for any of the integer programs \mathcal{F} given in either Definition 6 or Definition 7, and let $\text{SDPVal}(\mathcal{F})$ be the optimal value of the SDP relaxation for the instance \mathcal{F} . We have that $\text{SDPVal}(\mathcal{F}) \geq \text{OptVal}(\mathcal{F})$. Finally, let us define $\text{RndVal}(\mathcal{F})$ to be the expectation of the value of the objective function after randomized rounding procedure. The analysis of the approximation ratio for the algorithm boils down to proving $\text{RndVal}(\mathcal{F}) \geq \alpha \text{SDPVal}(\mathcal{F})$, where α is a constant that depends on the problem of interest. The way to calculate α is to look at the loss incurred by rounding at each constraint. Let us now show how this can be done for the CC-Max-Cut problem.

The expected value of each constraint $\frac{1-x_i x_j}{2}$ after rounding the SDP solution of CC-Max-Cut problem is $\frac{1-\mathbf{E}[\bar{y}_i \bar{y}_j]}{2}$, and therefore at each constraint the loss factor incurred by rounding is given as

$$\frac{1 - \mathbf{E}[\bar{y}_i \bar{y}_j]}{2} \cdot \frac{1}{(1 - \langle v_i, v_j \rangle)/2}.$$

Thus, in order to calculate the approximation ratio, we need to bound this expression from below. Let us first note that

$$\mathbf{E}[\bar{y}_1 \bar{y}_2] = 4\Gamma_{\bar{\rho}}\left(\frac{1-\mu_1}{2}, \frac{1-\mu_2}{2}\right) + \mu_1 + \mu_2 - 1,$$

where $\bar{\rho}$ is given as

$$\bar{\rho} = \frac{\rho - \mu_1 \mu_2}{\sqrt{1 - \mu_1^2} \sqrt{1 - \mu_2^2}}.$$

Then, the approximation ratio is lower bounded by the quantity α_{cut}^{cc} defined as the solution of the optimization problem

$$\alpha_{cut}^{cc} = \min_{(\mu_1, \mu_2, \rho) \in \mathbf{Conf}} \frac{2 - 4\Gamma_{\bar{\rho}}\left(\frac{1-\mu_1}{2}, \frac{1-\mu_2}{2}\right) - \mu_1 - \mu_2}{1 - \rho}.$$

² We assume that $\|w_i\| \neq 0$, since we can introduce a small perturbation to the values v_i without affecting the objective value too much.

Computing α_{cut}^{cc} is a hard global optimization problem, and therefore we resort to numerical computations to estimate it (we remark that the same approach is taken for a similar function in [23] and [3]). Extensive numerical experiments show that the minimum is attained at $\mu_1 = \mu_2 = \mu$, while the ρ is on the boundary of the polytope **Conf**, $\rho = -1 + 2|\mu|$. More precisely, the minimum is attained at $\mu \approx 0.27$, and $\rho \approx -0.575$, and it has a value of approximately 0.858.

Assuming that the minimum is attained at the configuration of the form $(\mu, \mu, -1+2\mu)$, $\mu > 0$, constant α_{cut}^{cc} can be found as the minimum of a function

$$\frac{1 - 2\Gamma_{\bar{\rho}}\left(\frac{1-\mu}{2}, \frac{1-\mu}{2}\right) - \mu}{1 - \mu},$$

where $\mu \in (0, 1)$. If we introduce $q = (1 - \mu)/2$, we can reexpress this function as

$$\alpha_{cut}^{cc}(q) = \frac{2q - 2\Gamma_{\bar{\rho}}(q)}{2q} = \frac{1 - \Gamma_{\bar{\rho}}(q) - \Gamma_{\bar{\rho}}(1 - q)}{2q}, \quad q \in (0, 1/2),$$

where in the last equality we used Lemma 5. Furthermore, $\bar{\rho} = -q/(1 - q)$. Similar analysis for CC-Max-2-Lin shows that the approximation ratio is the minimal value of the same function.

Straightforward calculations show that $\beta_{cut}^{cc}(q, -q/(1 - q))$ from Theorem 17 equals the value of $\alpha_{cut}^{cc}(q)$. Therefore, under the (mild) assumption that worst-case configurations indeed take the special form as explained above, our hardness result is sharp and the algorithm for CC-Max-Cut of Raghavendra and Tan is optimal on general instances of CC-Max-Cut / CC-Max-2-Lin.

In completely analogous way, we can conclude that the approximation ratio for CC-Max-2-Sat and Max- k -VC problems can be calculated as the minimum of the following function

$$\alpha_{2sat}^{cc}(q) = \frac{1 - \Gamma_{\bar{\rho}}(1 - q)}{2q}, \quad q \in (0, 1/2),$$

where $\bar{\rho} = -q/(1 - q)$. Numerical experiments show that $\alpha_{2sat}^{cc} \approx 0.929$, and that the minimum is attained at $q \approx 0.365$.

Again we have that the corresponding hardness expression from Theorem 18 satisfies $\beta_{vc}^{cc}(q, -q/(1 - q)) = \alpha_{2sat}^{cc}(q)$, implying (under the assumption on worst-case configurations) that the algorithm for CC-Max-2-Sat of Raghavendra and Tan is optimal.

5 Conclusion and Some Open Questions

We studied some of the cardinality constrained 2-CSPs, and assuming the Unique Games Conjecture derived hardness results which show that approximation ratios achieved by the algorithm described in [30] are optimal for CC-Max-2-Sat (and its special case Max- k -VC) and CC-Max-2-Lin (and its special case CC-Max-Cut). It would be interesting to derive UG-hardness for related CC-Max-CSPs of arity 2, most interestingly for the Max- k -DS problem. While super-constant hardness for Max- k -DS is currently known under the closely related Small-Set Expansion Hypothesis [29], it is not yet known whether the UGC implies hardness of Max- k -DS.

We also think it would be valuable to study whether we can achieve better approximation ratios or derive stronger hardness results for CC-Max-2-CSP with fixed values of the cardinality constraint q . Can the hardness curves of Theorem 17 and Theorem 18 depicted in Figure 1 be matched algorithmically for every q ?

Another interesting research direction would be to come up with hardness results for some other well-know Max-CSPs like Max-3-Sat, or even more ambitiously to extend the results of Raghavendra [28] and obtain tight hardness for all cardinality-constrained Max-CSPs.

References

- 1 Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof Verification and Hardness of Approximation Problems. In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*, pages 14–23, 1992. doi:10.1109/SFCS.1992.267823.
- 2 Sanjeev Arora and Shmuel Safra. Probabilistic Checking of Proofs; A New Characterization of NP. In *33rd Annual Symposium on Foundations of Computer Science, Pittsburgh, Pennsylvania, USA, 24-27 October 1992*, pages 2–13, 1992. doi:10.1109/SFCS.1992.267824.
- 3 Per Austrin. Balanced max 2-sat might not be the hardest. In *Proceedings of the 39th Annual ACM Symposium on Theory of Computing, San Diego, California, USA, June 11-13, 2007*, pages 189–197, 2007. doi:10.1145/1250790.1250818.
- 4 Per Austrin, Siavosh Benabbas, and Konstantinos Georgiou. Better Balance by Being Biased: A 0.8776-Approximation for Max Bisection. In *Proceedings of the Twenty-Fourth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2013, New Orleans, Louisiana, USA, January 6-8, 2013*, pages 277–294, 2013. doi:10.1137/1.9781611973105.21.
- 5 Per Austrin, Subhash Khot, and Muli Safra. Inapproximability of Vertex Cover and Independent Set in Bounded Degree Graphs. *Theory of Computing*, 7(1):27–43, 2011. doi:10.4086/toc.2011.v007a003.
- 6 Markus Bläser and Bodo Manthey. Improved Approximation Algorithms for Max-2SAT with Cardinality Constraint. In *Algorithms and Computation, 13th International Symposium, ISAAC 2002 Vancouver, BC, Canada, November 21-23, 2002, Proceedings*, pages 187–198, 2002. doi:10.1007/3-540-36136-7_17.
- 7 Joshua Brakensiek, Sivakanth Gopi, and Venkatesan Guruswami. CSPs with Global Modular Constraints: Algorithms and Hardness via Polynomial Representations. *Electronic Colloquium on Computational Complexity (ECCC)*, 26:13, 2019. URL: <https://ecc.ecc.wi.zmann.ac.il/report/2019/013>.
- 8 Andrei A. Bulatov. A Dichotomy Theorem for Nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 319–330, 2017. doi:10.1109/FOCS.2017.37.
- 9 Andrei A. Bulatov and Dániel Marx. The complexity of global cardinality constraints. *Logical Methods in Computer Science*, 6(4), 2010. doi:10.2168/LMCS-6(4:4)2010.
- 10 Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional Hardness for Approximate Coloring. *SIAM J. Comput.*, 39(3):843–873, 2009. doi:10.1137/07068062X.
- 11 Uriel Feige and Michael Langberg. The RPR² rounding technique for semidefinite programs. *J. Algorithms*, 60(1):1–23, 2006. doi:10.1016/j.jalgor.2004.11.003.
- 12 Alan M. Frieze and Mark Jerrum. Improved Approximation Algorithms for MAX k-CUT and MAX BISECTION. *Algorithmica*, 18(1):67–81, 1997. doi:10.1007/BF02523688.
- 13 Michel X. Goemans and David P. Williamson. .879-approximation algorithms for MAX CUT and MAX 2SAT. In *Proceedings of the Twenty-Sixth Annual ACM Symposium on Theory of Computing, 23-25 May 1994, Montréal, Québec, Canada*, pages 422–431, 1994. doi:10.1145/195058.195216.
- 14 Venkatesan Guruswami and Euiwoong Lee. Complexity of Approximating CSP with Balance / Hard Constraints. *Theory Comput. Syst.*, 59(1):76–98, 2016. doi:10.1007/s00224-015-9638-0.
- 15 Venkatesan Guruswami, Rajsekar Manokaran, and Prasad Raghavendra. Beating the Random Ordering is Hard: Inapproximability of Maximum Acyclic Subgraph. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008, Philadelphia, PA, USA*, pages 573–582, 2008. doi:10.1109/FOCS.2008.51.

- 16 Eran Halperin and Uri Zwick. A unified framework for obtaining improved approximation algorithms for maximum graph bisection problems. *Random Struct. Algorithms*, 20(3):382–402, 2002. doi:10.1002/rsa.10035.
- 17 Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. doi:10.1145/502090.502098.
- 18 Thomas Hofmeister. An Approximation Algorithm for MAX-2-SAT with Cardinality Constraint. In *Algorithms - ESA 2003, 11th Annual European Symposium, Budapest, Hungary, September 16-19, 2003, Proceedings*, pages 301–312, 2003. doi:10.1007/978-3-540-39658-1_29.
- 19 Howard J. Karloff and Uri Zwick. A 7/8-Approximation Algorithm for MAX 3SAT? In *38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997*, pages 406–415, 1997. doi:10.1109/SFCS.1997.646129.
- 20 Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings on 34th Annual ACM Symposium on Theory of Computing, May 19-21, 2002, Montréal, Québec, Canada*, pages 767–775, 2002. doi:10.1145/509907.510017.
- 21 Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O'Donnell. Optimal Inapproximability Results for MAX-CUT and Other 2-Variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007. doi:10.1137/S0097539705447372.
- 22 Subhash Khot and Oded Regev. Vertex Cover Might be Hard to Approximate to within $2-\epsilon$. In *18th Annual IEEE Conference on Computational Complexity (Complexity 2003), 7-10 July 2003, Aarhus, Denmark*, page 379, 2003. doi:10.1109/CCC.2003.1214437.
- 23 Michael Lewin, Dror Livnat, and Uri Zwick. Improved Rounding Techniques for the MAX 2-SAT and MAX DI-CUT Problems. In *Integer Programming and Combinatorial Optimization, 9th International IPCO Conference, Cambridge, MA, USA, May 27-29, 2002, Proceedings*, pages 67–82, 2002. doi:10.1007/3-540-47867-1_6.
- 24 Pasin Manurangsi. Almost-polynomial ratio ETH-hardness of approximating densest k-subgraph. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 954–961, 2017. doi:10.1145/3055399.3055412.
- 25 Pasin Manurangsi. A Note on Max k-Vertex Cover: Faster FPT-AS, Smaller Approximate Kernel and Improved Approximation. In *2nd Symposium on Simplicity in Algorithms, SOSA@SODA 2019, January 8-9, 2019 - San Diego, CA, USA*, pages 15:1–15:21, 2019. doi:10.4230/OASIcs.SOSA.2019.15.
- 26 Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. *Ann. of Math. (2)*, 171(1):295–341, 2010. doi:10.4007/annals.2010.171.295.
- 27 Ryan O'Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014. URL: <http://www.cambridge.org/de/academic/subjects/computer-science/algorithmics-complexity-computer-algebra-and-computational-g/analysis-boolean-functions>.
- 28 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*, pages 245–254, 2008. doi:10.1145/1374376.1374414.
- 29 Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In Leonard J. Schulman, editor, *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*, pages 755–764. ACM, 2010. doi:10.1145/1806689.1806792.
- 30 Prasad Raghavendra and Ning Tan. Approximating CSPs with global cardinality constraints using SDP hierarchies. In *Proceedings of the Twenty-Third Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2012, Kyoto, Japan, January 17-19, 2012*, pages 373–387, 2012. URL: <http://portal.acm.org/citation.cfm?id=2095149&CFID=63838676&CFTOKEN=79617016>, doi:10.1137/1.9781611973099.33.

- 31 Ran Raz. A Parallel Repetition Theorem. *SIAM J. Comput.*, 27(3):763–803, 1998. doi:10.1137/S0097539795280895.
- 32 Thomas J. Schaefer. The Complexity of Satisfiability Problems. In *Proceedings of the Tenth Annual ACM Symposium on Theory of Computing, STOC '78*, pages 216–226, New York, NY, USA, 1978. ACM. doi:10.1145/800133.804350.
- 33 Maxim Sviridenko. Best Possible Approximation Algorithm for MAX SAT with Cardinality Constraint. *Algorithmica*, 30(3):398–405, 2001. doi:10.1007/s00453-001-0019-5.
- 34 Yinyu Ye. A .699-approximation algorithm for Max-Bisection. *Math. Program.*, 90(1):101–111, 2001. doi:10.1007/PL00011415.
- 35 Dmitriy Zhuk. A Proof of CSP Dichotomy Conjecture. In *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 331–342, 2017. doi:10.1109/FOCS.2017.38.