Electronic Colloquium on Computational Complexity, Report No. 148 (2019)

Simultaneous Max-Cut is harder to approximate than Max-Cut

Amey Bhangale* Subhash Khot†

November 1, 2019

Abstract

A systematic study of simultaneous optimization of constraint satisfaction problems was initiated in [BKS15]. The simplest such problem is the simultaneous Max-Cut. [BKK+18] gave a .878-minimum approximation algorithm for simultaneous Max-Cut which is *almost optimal* assuming the Unique Games Conjecture (UGC). For a single instance Max-Cut, [GW95] gave an α_{GW}-approximation algorithm where $\alpha_{GW} \approx .87856720...$ which is *optimal* assuming the UGC.

It was left open whether one can achieve an α_{GW}-minimum approximation algorithm for simultaneous Max-Cut. We answer the question by showing that there exists an absolute constant $\varepsilon_0 \geqslant 10^{-5}$ such that it is NP-hard to get an $(\alpha_{GW} - \varepsilon_0)$-minimum approximation for simultaneous Max-Cut assuming the Unique Games Conjecture.

1 Introduction

Constraint satisfaction problems (CSPs) are among the most fundamental problems in computer science and Max-Cut is the most basic among those. In Max-Cut we are given an undirected (weighted) graph $G(V, E)$ on the vertex set V along with the edge set E. We assume that the total weight of edges is 1 and denote the number of vertices by n. The objective is to partition V into two sets $S, \overline{S}$ so as to maximize the total weight of *crossing* edges i.e. having one endpoint in S and the other in $\overline{S}$. Let us denote the cut value corresponding to the partition $(S, \overline{S})$ by $\mathbf{Cut}_G(S)$. Since Max-Cut is one of the classic NP-complete problems, we resort to finding an approximate solution. The seminal result of Goemans-Williamson [GW95] gave $\alpha_{GW} \approx .87856720...$ approximation algorithm

*Department of Computer Science and Engineering, University of California, Riverside. Part of this work was done while visiting the Simons Institute for the Theory of Computing, UC Berkeley. Email: ameyrbh@gmail.com

†Department of Computer Science, Courant Institute of Mathematical Sciences, New York University. Supported by the NSF Award CCF-1813438, the Simons Collaboration on Algorithms and Geometry, and the Simons Investigator Award. Email: khot@cs.nyu.edu

ISSN 1433-8092

for MAX-CUT. The exact value of the approximation factor is given by the following expression:

$$\alpha_{GW} := \min_{\rho \in [-1,0]} \frac{2\arccos(\rho)}{\pi(1-\rho)}.$$

In [BKS15], the authors initiate the study of simultaneous approximation algorithms for constraint satisfaction problems. In particular, the study of simultaneous MAX-CUT which we describe next and is the main focus of this paper. In simultaneous MAX-CUT the input consists of a collection of weighted undirected graphs $G_1, G_2, \ldots, G_k$ on the same set of vertices V but with different edge weights $E_1, E_2, \ldots, E_k$. The goal is to find a single cut $(S, \overline{S})$ which is *good* for each of G_i. The notion of how *good* the cut is needs to be defined formally. Following are the two notions that [BKS15] considered in their paper:

1. **Pareto approximation:** Suppose $(c_1, c_2, \ldots, c_k) \in [0,1]^k$ is such that there exists a partition $(S, \overline{S})$ such that $\mathbf{Cut}_{G_i}(S) \geqslant c_i$ for all $i \in [k]$. The objective is to find such a partition. An α*-Pareto approximation* algorithm in this context is a polynomial time algorithm, which when given $(c_1, c_2, \ldots, c_k) \in [0,1]^k$ as input, finds a partition $(S, \overline{S})$ such that $\mathbf{Cut}_{G_i}(S) \geqslant \alpha \cdot c_i$ for all $i \in [k]$.

2. **Minimum approximation:** This is the Pareto approximation problem when $c_1 = c_2 = \ldots = c_k$. Define the optimal value of the instance to be

$$c = \max_{S \subseteq V} \min_{i \in [k]} \mathbf{Cut}_{G_i}(S).$$

   An α*-minimum approximation* algorithm in this context is a polynomial time algorithm which finds a cut $(S, \overline{S})$ such that $\min_{i \in [k]} \mathbf{Cut}_{G_i}(S) \geqslant \alpha \cdot c$.

Note that an α-Pareto approximation gives an α-minimum approximation of simultaneous MAX-CUT. For any constant $k \geqslant 1$, [BKS15] gave $\frac{1}{2}$-Pareto approximation for simultaneous MAX-CUT which was improved to .878-Pareto approximation by [BKK$^+$18].

Theorem 1.1. *(Pareto approximation algorithm of [BKK$^+$18]) Given an instance $G_i(V, E_i)$ for $1 \leqslant i \leqslant k$ and $c_1.c_2.\ldots, c_k \in [0,1]$ with a guarantee that there exists a partition $(S^\star, \overline{S^\star})$ such that $\mathbf{Cut}_{G_i}(S^\star) \geqslant c_i$ for all i, there exists a randomized algorithm running in time $|V|^{poly(k)}$ which outputs a cut $(S, \overline{S})$ with a guarantee that $\mathbf{Cut}_{G_i}(S) \geqslant .878 \cdot c_i$ for all i.*

In terms of hardness of approximation, the Unique Games Conjecture by [Kho02] gives the tightness of the Goemans-Williamson algorithm for approximating MAX-CUT. [KKMO07] showed that if approximating a certain optimization problem called the *Unique Games* is NP-hard then it is NP-hard to approximate MAX-CUT better than α_{GW} factor. Trivially, the Unique Games Conjecture based hardness (UG-hard henceforth) of approximating MAX-CUT within a factor of $(\alpha_{GW} + \varepsilon)$ implies that getting an $(\alpha_{GW} + \varepsilon)$-Pareto approximation for simultaneous MAX-CUT is *also* UG-hard for all constants $\varepsilon > 0$. As

$.878 < \alpha_{GW}$, this leaves an intriguing question of achieving an α_{GW}-Pareto approximation for simultaneous MAX-CUT.

We answer this question in this paper by proving that there exists an absolute constant $\varepsilon_0 \geqslant 10^{-5}$ such that it is UG-hard to get an $(\alpha_{GW} - \varepsilon_0)$-minimum approximation (and hence $(\alpha_{GW} - \varepsilon_0)$-Pareto approximation) for simultaneous MAX-CUT, unlike the single instance MAX-CUT.

Theorem 1.2 (Main theorem)**.** *There exists an absolute constant $\varepsilon_0 \geqslant 10^{-5}$ such that assuming the Unique Games Conjecture, it is NP-hard to achieve $(\alpha_{GW} - \varepsilon_0)$-minimum approximation for simultaneous* MAX-CUT*.*

One interesting feature of our reduction is that the hard instance involves only *three* graphs! This should be compared with the algorithm of [BKK+18] form Theorem 1.1 which works for any constantly many number of instances of MAX-CUT. It will be interesting to know whether one can achieve α_{GW}-minimum approximation for the simultaneous MAX-CUT when the number of instances is *two*.

1.1 Organisation

We start with preliminaries in Section 2 where we formally define the simultaneous MAX-CUT problem, various distributions on the Boolean hypercube, invariance principle and the Unique Games Conjecture. In Section 3, we present the dictatorship tests for MAX-CUT and simultaneous MAX-CUT. Finally, in Section 4, we provide our reduction from the Unique Games to the simultaneous MAX-CUT.

2 Preliminaries

We first define the main problem that we study. Given an undirected weighted graph $G(V, E)$, the cut value of the partition $(S, \overline{S})$ of V, denoted by $\mathbf{Cut}_G(S)$, is defined to be the total weight of the edges whose endpoints are in different parts. The Max-Cut of a graph G is the maximum cut value over all the partitions of V.

Definition 2.1. *(Simultaneous* MAX-CUT*) An instance of simultaneous* MAX-CUT *is a collection of undirected weighted graphs $G_i(V, E_i)$, $1 \leqslant i \leqslant k$, on the same set of vertices.*

Given an instance $G_i(V, E_i)$, $1 \leqslant i \leqslant k$ of simultaneous MAX-CUT and $(c_1, c_2, \ldots, c_k) \in [0, 1]^k$ such that there exists a partition $(S, \overline{S})$ satisfying $\mathbf{Cut}_{G_i}(S) \geqslant c_i$ for all $i \in [k]$. The objective is to find such a partition. An α-Pareto approximation algorithm in this context is a polynomial time algorithm, which when given $(c_1, c_2, \ldots, c_k) \in [0, 1]^k$ as input, finds a partition $(S, \overline{S})$ such that $\mathbf{Cut}_{G_i}(S) \geqslant \alpha \cdot c_i$ for all $i \in [k]$.

We work with the problem of finding α-minimum approximation for simultaneous Max-Cut, which is a special case of the above problem. In this case, the optimum value is given by:

$$\textsc{Opt}(G_1, G_2, \ldots, G_k) := \max_{S \subseteq V} \min_{i \in [k]} \mathbf{Cut}_{G_i}(S).$$

An algorithm is called an α-minimum approximation for simultaneous Max-Cut if given input the graphs $G_1, G_2, \ldots, G_k$, it always outputs a cut $(T, \overline{T})$ such that

$$\min_{i \in [k]} \mathbf{Cut}_{G_i}(T) \geqslant \alpha \cdot \textsc{Opt}(G_1, G_2, \ldots, G_k).$$

For $a, b, c \in \mathbb{R}_{\geqslant 0}$ and a polynomial $P(x_1, x_2, \ldots, x_t)$, we define

$$\underset{x_1,\ldots,x_t \in [a,b]}{\texttt{range}} \{P(x_1, \ldots, x_t) \geqslant c\} := \{(x_1, \ldots, x_t) \mid x_i \in [a, b]\ \forall i \in [t] \text{ and } P(x_1, \ldots, x_t)) \geqslant c\}.$$

2.1 Analysis of Boolean functions

We will be working with functions $f : \{0,1\}^n \to \mathbb{R}$ on the Boolean hypercube. For $q \in [0,1]$, let μ_q be the distribution of a q-biased bit given as $\mu_q(1) = q$ and $\mu_q(0) = 1 - q$. Let $\mu_q^{\otimes n}$ be the corresponding product distribution on $\{0,1\}^n$. Let $L^2(\mu_q^{\otimes n})$ be the space of functions $f : \{0,1\}^n \to \mathbb{R}$ endowed with the distribution $\mu_q^{\otimes n}$. Also, let $\mu_q(f) := \mathbf{E}_{x \sim \mu_q^{\otimes n}}[f(x)]$.

Given x define the ρ-correlated copy y of x as follows:

Definition 2.2. *Given ρ and $x \sim \mu_q^{\otimes n}$ we write $y \sim N_\rho(x)$ to denote the ρ-correlated copy of x where the distribution $N_\rho(x)$ is as follows: Independently for each $i \in [n]$, if $x_i = 1$ then set $y_i = 1$ with probability $q + \rho(1-q)$, and $y_i = 0$ otherwise. If $x_i = 0$ then set $y_i = 1$ with probability $q - \rho q$, and $y_i = 0$ otherwise.*

Equivalently, first we set $y = x$, and then independently each coordinate of y is re-randomized w.r.t. μ_q with probability $(1 - \rho)$.

We will be interested in the setting when $\rho \leqslant 0$. In this case, if we want y to be distributed according to $\mu_q^{\otimes n}$ then ρ *cannot* be arbitrary in $[-1, 0]$. Specifically, for a given $q \in (0,1)$, ρ must be in the following interval:

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

As in [AS19], we will denote the above interval as $\kappa(q)$ for any given $q \in (0,1)$. Next we define the noise operator T_ρ over the probability space $L^2(\mu_q^{\otimes n})$.

Definition 2.3. *Let $q \in (0,1)$ and $\rho \in \kappa(q)$. The noise operator $T_\rho : L^2(\mu_q^{\otimes n}) \to L^2(\mu_q^{\otimes n})$ is given as follows:*

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

Definition 2.4 (Influence)**.** *Let $f \in L^2(\mu_q^{\otimes n})$. The influence of the i^{th} variable on f, denoted by $\mathbf{Inf}_i(f)$ is defined as:*

$$\mathbf{Inf}_i(f) = \mathop{\mathbf{E}}_{x \sim \mu_q^{\otimes n}} [\mathbf{Var}_{x_i \sim \mu_q}[f(x) | x_1, x_2, \ldots, x_{i-1}, x_{i+1}, \ldots, x_n]].$$

The useful property of the operator T_ρ is that if $\mathbf{Var}[f]$ is bounded then its image under T_ρ has a bounded number of influential variables.

Lemma 2.5. *Let $q \in (0,1)$ and $\rho \in \kappa(q)$ and $f \in L^2(\mu_q^{\otimes n})$. Then, for any $\tau > 0$ we have*

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

We have the following definition for functions whose all the influences are low (under the map T_ρ).

Definition 2.6. *Let $q \in (0,1)$ and $0 < \varepsilon, \delta < 1$. A function $f \in L^2(\mu_q^{\otimes n})$ is called (ε, δ)-quasirandom if for all $i \in [n]$, we have $\mathbf{Inf}_i[T_{1-\delta} f] \leqslant \varepsilon$.*

2.2 Invariance Principle

We need the following definition related to correlated spaces defined by Mossel [Mos10].

Definition 2.7. *Let $(\Omega_1 \times \Omega_2, \mu)$ be a finite correlated space, the correlation between Ω_1 and Ω_2 with respect to μ us defined as*

$$\rho(\Omega_1, \Omega_2; \mu) := \sup_{\substack{f:\Omega_1 \to \mathbb{R}, g:\Omega_2 \to \mathbb{R}, \\ \mathbf{Var}[f] = \mathbf{Var}[g] = 1}} \mathbf{Cov}[f, g].$$

We will need the following Gaussian stability measure in our analysis:

Definition 2.8. *Let $\phi : \mathbb{R} \to [0,1]$ be the cumulative distribution function of the standard Gaussian random variable. For a parameter $\rho, \nu_1, \nu_2 \in [0,1]$, we define the following two quantities:*

$$\underline{\Gamma}_\rho(\nu_1, \nu_2) = \Pr[X \leqslant \phi^{-1}(\nu_1), Y \geqslant \phi^{-1}(1 - \nu_2)]$$

$$\overline{\Gamma}_\rho(\nu_1, \nu_2) = \Pr[X \leqslant \phi^{-1}(\nu_1), Y \leqslant \phi^{-1}(\nu_2)]$$

where X and Y are two standard Gaussian variables with covariance ρ. We also define $\Gamma_\rho(\nu) = \overline{\Gamma}_\rho(\nu, \nu)$.

We are now ready to state a version of invariance principle from [Mos10] which follows from Theorem 3.1 in [DMR09] that we need for our reduction. For variables $\varepsilon_1, \varepsilon_2, \varepsilon_3, \ldots$, by $\varepsilon_1(\varepsilon_2, \varepsilon_3, \ldots)$ we mean ε_1 is a function of $\varepsilon_2, \varepsilon_3, \ldots$ such that $\varepsilon_1 \to 0$ as all $\varepsilon_2, \varepsilon_3, \ldots \to 0$.

Theorem 2.9 ([Mos10, DMR09]). *Let $(\Omega_1 \times \Omega_2, \mu)$ be a finite correlated space, the correlation between Ω_1 and Ω_2 with respect to μ is $\rho \in [0,1]$. Then for any $\tau > 0$ there exists $\varepsilon(\tau) > 0, \delta(\tau) > 0$ such that if $f : \Omega_1^n \to [0,1]$ and $g : \Omega_2^n \to [0,1]$ are two functions satisfying*

$$\min(\mathbf{Inf}_i(T_{1-\delta} f), \mathbf{Inf}_i(T_{1-\delta} g)) \leqslant \varepsilon \tag{1}$$

for all $i \in [n]$, then it holds that

$$\underline{\Gamma}_\rho(\nu_1, \nu_2) - \tau \leqslant \mathop{\mathbf{E}}_{(x,y)\sim\mu^{\otimes n}}[f(x)g(y)] \leqslant \overline{\Gamma}_\rho(\nu_1, \nu_2) + \tau$$

where $\nu_1 = \mathbf{E}[f]$, $\nu_2 = \mathbf{E}[g]$.

Remark 2.10. *One difference between the versions of invariance principle in Mossel [Mos10] and Dinur et al. [DMR09] is that in [Mos10] instead of a* min *in (1), it was a* max. *This improvement was crucial for hardness of graph coloring in [DMR09]. For our hardness result, the difference is not important.*

We will be working with correlated spaces with negative correlation. The following corollary follows from the above theorem.

Corollary 2.11. *Assume the settings in Theorem 2.9 except $\rho \in [-1, 0)$ then it holds that*

$$\overline{\Gamma}_\rho(\nu_1, \nu_2) - \tau \leqslant \mathop{\mathbf{E}}_{(x,y)\sim\mu^{\otimes n}}[f(x)g(y)].$$

Proof. Define $f'(x) = 1 - f(-x)$ and let $\rho' = -\rho$. We apply Theorem 2.9 to f', g and ρ'

$$\begin{aligned}
\mathbf{E}[f(x)g(y)] &= \mathbf{E}[g(y)] - \mathbf{E}[f'(-x)g(y)] \\
&\geqslant \nu_2 - \overline{\Gamma}_{\rho'}(1-\nu_1, \nu_2) - \tau \\
&= \nu_2 - \overline{\Gamma}_{\rho'}(1-\nu_1, \nu_2) - \underline{\Gamma}_{\rho'}(\nu_1, \nu_2) + \underline{\Gamma}_{\rho'}(\nu_1, \nu_2) - \tau
\end{aligned}$$

Now, $\overline{\Gamma}_{\rho'}(1-\nu_1, \nu_2) + \underline{\Gamma}_{\rho'}(\nu_1, \nu_2) = \overline{\Gamma}_{\rho'}(\nu_2, 1-\nu_1) + \underline{\Gamma}_{\rho'}(\nu_2, \nu_1) = \nu_2$. Therefore,

$$\begin{aligned}
\mathbf{E}[f(x)g(y)] &\geqslant \underline{\Gamma}_{\rho'}(\nu_1, \nu_2) - \tau \\
&= \overline{\Gamma}_\rho(\nu_1, \nu_2) - \tau
\end{aligned}$$

□

2.3 Unique Games

Our hardness result is based on the Unique Games Conjecture. First, we define what the Unique Game is:

Definition 2.12 (Unique Games)**.** *An instance $G = (U, V, E, [L], \{\pi_e\}_{e\in E})$ of the Unique Games constraint satisfaction problem consists of a bi-regular bipartite graph (U, V, E), an alphabet $[L]$ and a permutation map $\pi_e : [L] \to [L]$ for every edge $e \in E$. Given a labeling $\ell : U \cup V \to [L]$, , an edge $e = (u, v)$ is said to be satisfied by ℓ if $\pi_e(\ell(v)) = \ell(u)$.*

G is said to be at most δ-satisfiable *if every labeling satisfies at most a δ fraction of the edges.*

The following is a conjecture by Khot [Kho02] which has been used to prove many *tight* inapproximability results.

Conjecture 2.13 (Unique Games Conjecture[Kho02])**.** *For every sufficiently small $\delta > 0$ there exists $L \in \mathbb{N}$ such that the following holds. Given a an instance $(U, V, E, [L], \{\pi_e\}_{e\in E})$ of Unique Games it is NP-hard to distinguish between the following two cases:*

- *YES case: There exist an assignment that satisfies at least $(1 - \delta)$ fraction of the edges.*
- *NO case: Every assignment satisfies at most δ fraction of the edge constraints.*

3 Dictatorship Tests

A function $f : \{0, 1\}^n \to \mathbb{R}$ is called a dictator function if $f(x_1, x_2, \ldots, x_n) = x_i$ for some $i \in [n]$. Dictatorship tests are designed to distinguish between the cases when f is a dictator function and f is an (ε, δ)-quasirandom function for small enough $\varepsilon, \delta > 0$.

3.1 Dictatorship Test for Max-Cut

The α_{GW} Unique Games hardness of Max-Cut relies on the analysis of a certain dictatorship test that we describe next. This will lead us to our dictatorship test for simultaneous Max-Cut. Consider the following test:

Given $f : \{0, 1\}^n \to \{0, 1\}$

1. Select $x \in \{0, 1\}^n$ uniformly at random.
2. Select a ρ-correlated copy y of x i.e. independently for each $i \in [n]$ set $y_i = x_i$ w.p. $\frac{1+\rho}{2}$ and set $y_i = \overline{x_i}$ w.p. $\frac{1-\rho}{2}$.
3. Check if $f(x) \neq f(y)$.

We have the following completeness property of the dictatorship test, which is easy to show.

Lemma 3.1. *If f is a dictator function, then the test passes with probability $\frac{1-\rho}{2}$.*

The following soundness of the test relies on the "Majority of the Stablest" theorem, which roughly states that among all the Boolean functions with all the influences low, Majority function is the most stable under 'positive' perturbation.

Lemma 3.2 ([MOO05])**.** *For $\rho \in [-1, 0)$, if f is (ε, δ)-quasirandom, then the test passes with probability at most $\frac{\arccos(\rho)}{\pi} + \tau(\varepsilon, \delta)$.*

This dictatorship test can be *composed* with Unique Games [KKMO07] which gives α_{GW}-hardness of approximation for Max-Cut, where α_{GW} is given by the following expression.

$$\min_{\rho \in [-1,0)} \frac{\frac{\arccos(\rho)}{\pi}}{\frac{1-\rho}{2}} = \alpha_{GW} = .87856720...$$

3.2 Dictatorship Test for simultaneous Max-Cut

In the above dictatorship test, we get a family of graphs parameterized by the quantity ρ. This might give a way to construct multiple instances of Max-Cut, one for each $\rho \in (-1, 1)$. However, this will not work and instead we will construct instances whose vertex set is concentrated around the $q \cdot n^{th}$ slice of the hypercube for some $q \in (0, 1)$. This will give us the family of graphs for each $q \in (0, 1)$ and ρ.

Our final dictatorship test for the simultaneous Max-Cut problem will consist of *three* graphs, G_1 on the qn^{th} slice, G_2 on the $(1-q)n^{th}$ slice and G_3 will be a bipartite graph between the $qn^{th}$and $(1-q)n^{th}$ slice of the Boolean hypercube $\{0, 1\}^n$.

Definition 3.3. *(ρ-correlated μ_q strings) For every $q \in [0, 1]$ and $\rho \in \kappa(q)$, define $\mathcal{A}_{\rho,q}^{\otimes n}$ to be the product distribution on $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$ where, $\mathcal{A}_{\rho,q} : \{0, 1\}^2 \to \mathbb{R}_{\geqslant 0}$ is defined as follows:*

$$\begin{aligned}
\mathcal{A}_{\rho,q}(0, 0) &= (1 - q) - t \\
\mathcal{A}_{\rho,q}(0, 1) &= t \\
\mathcal{A}_{\rho,q}(1, 0) &= t \\
\mathcal{A}_{\rho,q}(1, 1) &= q - t,
\end{aligned}$$

where $t = (q - q^2)(1 - \rho)$. As mentioned before, ρ in the above definition must satisfy the following property

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

Definition 3.4. *(ρ-correlated (x, y) where $x \sim \mu_q^{\otimes n}$ and $y \sim \mu_{(1-q)}^{\otimes n}$) For every $q \in [0, 1]$ and $\rho \in [-1, 1]$, define $\mathcal{B}_{\rho,q}^{\otimes n}$ to be the product distribution on $(x, y) \in \{0, 1\}^n \times \{0, 1\}^n$*

where, $\mathcal{B}_{\rho,q} : \{0,1\}^2 \to \mathbb{R}_{\geqslant 0}$ *is defined as follows:*

$$\begin{aligned}
\mathcal{B}_{\rho,q}(0,0) &= t \\
\mathcal{B}_{\rho,q}(0,1) &= (1-q) - t \\
\mathcal{B}_{\rho,q}(1,0) &= q - t \\
\mathcal{B}_{\rho,q}(1,1) &= t,
\end{aligned}$$

where $t = (q - q^2)(1+\rho)$*. Note that* ρ *in the above definition must satisfy the following property:*

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

We will define a simultaneous Max-Cut instance on the vertex set $\{0,1\}^n$. The instance consists of three weighted graphs G_1, G_2 and G_3. We fix $q_\star = .58$, $\rho_1 = -\frac{1-q_\star}{q_\star}$ and $\rho_2 = \frac{2q_\star^2-1}{2q_\star(1-q_\star)}$.

Graph G_1: G_1 is concentrated on the $q_\star n^{th}$ slice of the hypercube. More formally, the edge distribution of this graph is given by the distribution $\mathcal{A}_{\rho_1,q_\star}^{\otimes n}$.

Graph G_2: G_2 is concentrated on the $(1-q_\star)n^{th}$ slice of the hypercube. Formally, the edge distribution of this graph is given by the distribution $\mathcal{A}_{\rho_1,(1-q_\star)}^{\otimes n}$.

Graph G_3: This is roughly a bipartite graph between the $q_\star n^{th}$ and $(1-q_\star)n^{th}$ slices of the hypercube. The edge distribution is given by the distribution $\mathcal{B}_{\rho_2,q_\star}^{\otimes n}$.

A few remarks about the choice of parameters: We arrive at the choice of $q_\star = .58$ by doing numerical calculations. Setting $\rho_1 = -\frac{1-q_\star}{q_\star}$ is a natural choice as it is the maximum negative correlation that the two $q_\star$-biased bits can have. Finally, $\rho_2 = \frac{2q_\star^2-1}{2q_\star(1-q_\star)}$ is chosen such that the following is satisfied:

$$\Pr_{(x_i,y_i)\sim\mathcal{A}_{\rho_1,q_\star}}[x_i \neq y_i] = \Pr_{(x_i,y_i)\sim\mathcal{B}_{\rho_2,q_\star}}[x_i \neq y_i].$$

3.2.1 Completeness

Lemma 3.5. *If* f *is a dictator function then the value of the cut induced by* f *is* $2(1-q_\star)$ *for all* G_1, G_2, G_3.

Proof. The proof is easy in this case. Suppose f is an i^{th} dictator for some $i \in [n]$. This induces a cut $(S_f, \overline{S}_f)$ where $S_f = \{x \in \{0,1\}^n \,|\, x_i = 0\}$. In this case, $\mathbf{Cut}_{G_1}(S_f)$ is equal to the probability that (x_i, y_i) sampled from $\mathcal{A}_{\rho_1, q_\star}$ are not equal. This is precisely $2(q_\star - q_\star^2)(1-\rho_1)$ which is equal to $2(1-q_\star)$ by the choice of $\rho_1 = -\frac{1-q_\star}{q_\star}$.

Similarly, $\mathbf{Cut}_{G_2}(S_f)$ is equal to the probability that (x_i, y_i) sampled from $\mathcal{A}_{\rho_1,(1-q_\star)}$ are not equal. This is also $2(1-q_\star)$.

For G_3,

$$\mathbf{Cut}_{G_3}(S_f) = \Pr_{(x_i,y_i)\sim\mathcal{B}_{\rho_2,q_\star}}[x_i \neq y_i] = 1 - 2(q_\star - q_\star^2)(1+\rho_2).$$

By the choice of ρ_2 this also equals to $2(1-q_\star)$. □

3.2.2 Soundness

Lemma 3.6. *Let $f : \{0,1\}^n \to \{0,1\}$ be an (ε, δ)-quasirandom function and let $(S_f, \overline{S}_f)$ be the cut induced by f. Then*

$$\min_{i\in[3]} \mathbf{Cut}_{G_i}(S_f) \leqslant (\alpha_{GW} - 10^{-5}) \cdot 2(1-q_\star) + \tau(\varepsilon, \delta).$$

Proof. The proof is as follows:

1. We have an (ε, δ)-quasirandom function $f : \{0,1\}^n \to \{0,1\}$. Invariance principle says that in order to get at least $(\alpha_{GW} - 10^{-5})$ approximation for G_1, the density of function $\mu_{q_\star}(f)$ must be in some range. This essentially follows from the analysis of Austrin et al. [AKS11, AS19]. Furthermore, the invariance principle precisely tells us that this is similar to what approximation ratio the *biased hyperplane rounding* algorithm of [BKK+18] gives us on a pair of vectors with SDP biases $q_\star$ when rounded using rounding bias $\mu_{q_\star}(f)$. (See [BKK+18] for the formal definitions of SDP bias and rounding bias). More formally, if the $\mu_{q_\star}(f) = \nu_1$ then the cut value is bounded as follows:

$$\begin{aligned}
\mathbf{Cut}_{G_1}(S_f) &= \mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes n}}\left[\frac{1-(1-2f(x))(1-2f(y))}{2}\right] \\
&= \mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes n}}[f(x)+f(y)-2f(x)f(y)] \\
&= \nu_1 + \nu_1 - 2\mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes n}}[f(x)f(y)] \\
&\leqslant 2\nu_1 - 2\overline{\Gamma}_{\rho_1}(\nu_1) + \tau_1(\varepsilon,\delta),
\end{aligned}$$

   where the last inequality follows from Corollary 2.11. Let us define the following range:

$$R_1(\varepsilon,\delta) := \mathop{\mathtt{range}}_{\nu_1\in[0,1]}\left\{\frac{2\nu_1 - 2\Gamma_{\rho_1}(\nu_1) + \tau_1(\varepsilon,\delta)}{2(1-q_\star)} \geqslant (\alpha_{GW} - 10^{-5})\right\}.$$

$R_1(\varepsilon, \delta)$ is the set of all biases $\mu_{q_\star}(f)$ that gives $\mathbf{Cut}_{G_1}(S_f)$ which is at least $(\alpha_{GW} - 10^{-5})$ factor greater than $2(1-q_\star)$. For a sufficiently small $\varepsilon, \delta > 0$ and our given values of $q_\star$ and ρ_1, numerical calculations show that

$$R_1(\varepsilon, \delta) \subseteq [.43676765, .56323235].$$

2. Same is true for G_2. More formally, if the $\mu_{1-q_\star}$ measure of f is ν_2 then the cut value is bounded above by $2\nu_2 - 2\Gamma_{\rho_1}(\nu_2)$ and we have

$$R_2(\varepsilon, \delta) := \underset{\nu_2 \in [0,1]}{\texttt{range}} \left\{ \frac{2\nu_2 - 2\Gamma_{\rho_1}(\nu_2) + \tau_2(\varepsilon, \delta)}{2(1-q_\star)} \geqslant (\alpha_{GW} - 10^{-5}) \right\}.$$

3. This fixes possible densities of f with respect to the $\mu_{q_\star}^{\otimes n}$ and $\mu_{(1-q_\star)}^{\otimes n}$ distributions. Both these densities should lie in $[.43676765, .56323235]$ if we want $\mathbf{Cut}_{G_1}(S_f) \geqslant (\alpha_{GW} - 10^{-5}) \cdot 2(1-q_\star)$ and $\mathbf{Cut}_{G_2}(S_f) \geqslant (\alpha_{GW} - 10^{-5}) \cdot 2(1-q_\star)$. Now we use the full power of the invariance principle to claim that the value of the cut given by such an f is similar to what the biased hyperplane rounding gives us on the graph G_3.

$$\begin{aligned}
\mathbf{Cut}_{G_3}(S_f) &= \underset{(x,y)\sim \mathcal{B}_{\rho_2, q_\star}^{\otimes n}}{\mathbf{E}} \left[\frac{1 - (1-2f(x))(1-2f(y))}{2} \right] \\
&= \underset{(x,y)\sim \mathcal{B}_{\rho_2, q_\star}^{\otimes n}}{\mathbf{E}} [f(x) + f(y) - 2f(x)f(y)] \\
&= \nu_1 + \nu_2 - 2 \underset{(x,y)\sim \mathcal{B}_{\rho_2, q_\star}^{\otimes n}}{\mathbf{E}} [f(x)f(y)] \\
&\leqslant \nu_1 + \nu_2 - \overline{\Gamma}_{\rho_2}(\nu_1, \nu_2) + \tau_3(\varepsilon, \delta).
\end{aligned}$$

Here again, the last inequality follows from Corollary 2.11. By doing numerical calculations, we show that for the following range

$$R(\varepsilon, \delta) := \underset{\nu_1, \nu_2 \in [0,1]}{\texttt{range}} \left\{ \frac{\nu_1 + \nu_2 - 2\overline{\Gamma}_{\rho_2}(\nu_1, \nu_2) + \tau_3(\varepsilon, \delta)}{2(1-q_\star)} \geqslant (\alpha_{GW} - 10^{-5}) \right\},$$

$R(\varepsilon, \delta) \cap (R_1(\varepsilon, \delta) \times R_2(\varepsilon, \delta)) = \emptyset$ for sufficiently small $\varepsilon, \delta > 0$.

Therefore, no matter which f we start with, if it is (ε, δ)-quasirandom for sufficiently small $\varepsilon, \delta > 0$, then there exists an $i \in [3]$ such that the cut guaranteed by S_f on G_i is strictly less that $(\alpha_{GW} - 10^{-5}) \cdot 2(1-q_\star) + \tau(\varepsilon, \delta)$. □

4 Actual Reduction

In this section we give a reduction from Unique Games to the simultaneous Max-Cut problem. Given an instance $G = (U, V, E, [L], \{\pi_e\}_{e\in E})$ of the Unique Games, we reduce it to a simultaneous Max-Cut instance $\mathcal{I}$ on the vertex set $\mathcal{V} = V \times 2^{[L]} = \{(v, x) \mid v \in V, x \in \{0,1\}^L\}$.

The instance will involve three weighted graphs $\mathcal{G}_1(\mathcal{V}, \mathcal{E}_1), \mathcal{G}_2(\mathcal{V}, \mathcal{E}_2)$ and $\mathcal{G}_3(\mathcal{V}, \mathcal{E}_3)$ on the common vertex set $\mathcal{V}$. We fix the following parameters: $q_\star = .58, \rho_1 = -\frac{1-q_\star}{q_\star}$ and $\rho_2 = \frac{2q_\star^2-1}{2q_\star(1-q_\star)}$. For a string $x \in \{0,1\}^L$ and a permutation $\pi : [L] \to [L]$, define $x \circ \pi \in \{0,1\}^L$ such that $(x \circ \pi)_i = x_{\pi(i)}$ for all $i \in [L]$. The respective edge weights are given by the following distributions:

1. $\mathcal{E}_1$: Select $u \in U$ uniformly at random and $v_1, v_2 \sim N(u)$ independently and uniformly at random. Select (x, y) according to $\mathcal{A}_{\rho_1,q_\star}^{\otimes L}$ and output $(v_1, x \circ \pi_{uv_1}^{-1}), (v_2, y \circ \pi_{uv_2}^{-1})$.
2. $\mathcal{E}_2$: Select $u \in U$ uniformly at random and $v_1, v_2 \sim N(u)$ independently and uniformly at random. Select (x, y) according to $\mathcal{A}_{\rho_1,(1-q_\star)}^{\otimes L}$ and output $(v_1, x\circ\pi_{uv_1}^{-1}), (v_2, y\circ\pi_{uv_2}^{-1})$.
3. $\mathcal{E}_3$: Select $u \in U$ uniformly at random and $v_1, v_2 \sim N(u)$ independently and uniformly at random. Select (x, y) according to $\mathcal{B}_{\rho_2,q_\star}^{\otimes L}$ and output $(v_1, x \circ \pi_{uv_1}^{-1}), (v_2, y \circ \pi_{uv_2}^{-1})$.

We now prove the completeness and the soundness of the reduction.

Lemma 4.1. *(Completeness) If the Unique Games instance G is $(1 - \frac{\eta}{2})$-satisfiable then there exists a cut $(\mathcal{S}, \overline{\mathcal{S}})$ such that*

$$\min_{i\in[3]} \mathbf{Cut}_{\mathcal{G}_i}(\mathcal{S}) \geqslant 2(1 - q_\star) - \eta.$$

Lemma 4.2. *(Soundness) There exist absolute constants $\varepsilon_0 \geqslant 10^{-5}$ and $0 < \eta_0 < 1$ such that for all $0 < \eta \leqslant \eta_0$ and $\varepsilon(\eta/2), \delta(\eta/2)$ from Theorem 2.9, if there exists a cut $(S, \overline{S})$ such that*

$$\min_{i\in[3]} \mathbf{Cut}_{\mathcal{G}_i}(S) \geqslant (\alpha_{GW} - \varepsilon_0)(2(1 - q_\star) - \eta),$$

then there exists an assignment to the Unique Games instance G which satisfies at least $\eta' = \eta \cdot \frac{\varepsilon}{2} \cdot \frac{\varepsilon\cdot e\cdot \ln(1-\delta)}{2}$ fraction of the constraints.

The above two lemmas along with Conjecture 2.13 show that assuming the Unique Games Conjecture, it is NP-hard to get an α-minimum approximation for simultaneous Max-Cut where $\alpha \leqslant \alpha_{GW} - 10^{-5}$. This proves Theorem 1.2. We now prove the completeness and soundness of the reduction.

Proof of Lemma 4.1. Let $\sigma : U \cup V \to [L]$ be an assignment to the Unique Games instance G which satisfies at least $(1-\eta)$ fraction of the constraints. Consider the following partition $(\mathcal{S}, \overline{\mathcal{S}})$ of $\mathcal{V}$ where

$$\mathcal{S} = \{(v, x) \mid v \in V, x_{\sigma(v)} = 0\}.$$

Let us analyze the value of this cut for the graph $\mathcal{G}_1$:

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_1}(\mathcal{S}) &= \mathop{\mathbf{E}}_{u \in U} \mathop{\mathbf{E}}_{v_1, v_2 \in N(u)} \Pr_{(x,y) \sim \mathcal{A}_{\rho_1, q_\star}^{\otimes L}} [(v_1, x \circ \pi_{uv_1}^{-1}), (v_2, y \circ \pi_{uv_2}^{-1}) \text{ in different parts}] \\
&= \mathop{\mathbf{E}}_{u \in U} \mathop{\mathbf{E}}_{v_1, v_2 \in N(u)} \Pr_{(x,y) \sim \mathcal{A}_{\rho_1, q_\star}^{\otimes L}} [((x \circ \pi_{uv_1}^{-1}))_{\sigma(v_1)} \neq (y \circ \pi_{uv_2}^{-1})_{\sigma(v_2)}] \\
&= \mathop{\mathbf{E}}_{u \in U} \mathop{\mathbf{E}}_{v_1, v_2 \in N(u)} \Pr_{(x,y) \sim \mathcal{A}_{\rho_1, q_\star}^{\otimes L}} [x_{\pi_{uv_1}^{-1}(\sigma(v_1))} \neq y_{\pi_{uv_2}^{-1}(\sigma(v_2))}] \\
&\geqslant (1-\eta) \Pr_{(x,y) \sim \mathcal{A}_{\rho_1, q_\star}^{\otimes L}} [x_{\sigma(u)} \neq y_{\sigma(u)}] \\
&= (1-\eta) \cdot 2(q_\star - q_\star^2)(1-\rho_1) \\
&= (1-\eta) \cdot 2(1-q_\star) \\
&\geqslant 2(1-q_\star) - \eta
\end{aligned}$$

where the first inequality uses the fact that with probability at least $1-\eta$, both the constraints on the edges (u, v_1) and (u, v_2) are satisfied by the assignment σ. Using similar calculations, we can show that

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_2}(\mathcal{S}) &\geqslant (1-\eta) \cdot 2(q_\star - q_\star^2)(1-\rho_1) \geqslant 2(1-q_\star) - \eta \\
\mathbf{Cut}_{\mathcal{G}_3}(\mathcal{S}) &\geqslant (1-\eta) \cdot (1 - 2(q_\star - q_\star^2)(1+\rho_2)) \geqslant 2(1-q_\star) - \eta.
\end{aligned}$$

Thus, we have

$$\min_{i \in [3]} \mathbf{Cut}_{\mathcal{G}_i}(\mathcal{S}) \geqslant 2(1-q_\star) - \eta.$$

□

We now prove the main soundness lemma:

Proof of Lemma 4.2. Suppose the value of the Unique Games instance is at most η'. Let $f : V \times 2^{[L]} \to \{0,1\}$ be the indicator function of the cut $(\mathcal{S}, \overline{\mathcal{S}})$. We will show that

$$\min_{i \in [3]} \mathbf{Cut}_{\mathcal{G}_i}(S) \leqslant (\alpha_{GW} - \varepsilon_0)(2(1-q_\star) - \eta).$$

We start with analysing the value $\mathbf{Cut}_{\mathcal{G}_1}(\mathcal{S})$:

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_1}(\mathcal{S}) &= \mathop{\mathbf{E}}_{u\in U}\ \mathop{\mathbf{E}}_{v_1,v_2\in N(u)}\ \Pr_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes L}}\left[f(v_1, x\circ\pi_{uv_1}^{-1}) \neq f(v_2, y\circ\pi_{uv_2}^{-1})\right] \\
&= \mathop{\mathbf{E}}_{u\in U}\ \mathop{\mathbf{E}}_{v_1,v_2\in N(u)}\ \mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes L}}\left[\frac{1}{2} - \frac{(1-2f(v_1, x\circ\pi_{uv_1}^{-1}))(1-2f(v_2, y\circ\pi_{uv_2}^{-1}))}{2}\right] \\
&= \mathop{\mathbf{E}}_{u\in U}\ \mathop{\mathbf{E}}_{v_1,v_2\in N(u)}\ \mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes L}}\left[f(v_1, x\circ\pi_{uv_1}^{-1}) + f(v_2, y\circ\pi_{uv_2}^{-1}) - 2f(v_1, x\circ\pi_{uv_1}^{-1})f(v_2, y\circ\pi_{uv_2}^{-1})\right]
\end{aligned}$$

Define $f_v(x) := f(v,x)$ for $v \in V$ and $f_u(x) := \mathbf{E}_{v\sim N(u)}\left[f_v(x\circ\pi_{uv}^{-1})\right]$ for $u \in U$. Let $\nu_q^u(f) = \mathbf{E}_{x\sim\mu_q^{\otimes L}}[f_u(x)]$ be the q-biased measure of the function f_u and $\nu_q(f) = \mathbf{E}_{u\in U}[\nu_q^u(f)]$ be the average q-biased measure of f. Since we sample $v_1, v_2 \in N(u)$ independently, we have

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_1}(\mathcal{S}) &= \mathop{\mathbf{E}}_{u\in U}\ \mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes L}}\left[f_u(x) + f_u(y) - 2f_u(x)f_u(y)\right] \\
&= 2\cdot\nu_{q_\star}(f) - 2\mathop{\mathbf{E}}_{u\in U}\ \mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes L}}\left[f_u(x)f_u(y)\right].
\end{aligned}$$

We now show that the expectation in the above expression is lower bounded by the quantity $\overline{\Gamma}_{\rho_1}(\nu_{q_\star}^u(f), \nu_{q_\star}^u(f)) - \frac{\eta'}{2}$ unless the value of the Unique Games instance is at least η'.

Claim 4.3. *For at least $(1-\eta)$ fraction of $u \in U$,*

$$\mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes L}}\left[f_u(x)f_u(y)\right] \geqslant \overline{\Gamma}_{\rho_1}(\nu_{q_\star}^u(f), \nu_{q_\star}^u(f)) - \frac{\eta}{2}.$$

Proof. Consider $f_u \in L^2(\mu_{q_\star}^{\otimes n})$ and suppose the claim is not true and we have for at least η fraction of $u \in U$,

$$\mathop{\mathbf{E}}_{(x,y)\sim\mathcal{A}_{\rho_1,q_\star}^{\otimes L}}\left[f_u(x)f_u(y)\right] \leqslant \overline{\Gamma}_{\rho_1}(\nu_{q_\star}^u(f), \nu_{q_\star}^u(f)) - \frac{\eta}{2}.$$

Then using Corollary 2.11, there exists $\varepsilon(\eta/2), \delta(\eta/2) > 0$ such that for at least η fraction of f_u, we have that $\mathbf{Inf}_i(T_{1-\delta}f_u) \geqslant \varepsilon$ for some $i \in [L]$. Since $f_u(x) := \mathbf{E}_{v\sim N(u)}\left[f_v(x\circ\pi_{uv}^{-1})\right]$ and $\mathbf{Inf}_i$ is a convex function, we have

$$\begin{aligned}
\mathop{\mathbf{E}}_{v\sim N(u)}\left[\mathbf{Inf}_i(T_{1-\delta}(f_v(x\circ\pi_{uv}^{-1})))\right] \geqslant \varepsilon \implies& \\
\mathop{\mathbf{E}}_{v\sim N(u)}\left[\mathbf{Inf}_{\pi_{uv}(i)}(T_{1-\delta}f_v)\right] \geqslant \varepsilon.&
\end{aligned}$$

Thus, if $\mathbf{Inf}_i(T_{1-\delta} f_u) \geqslant \varepsilon$, then by an averaging argument, for at least $\varepsilon/2$ fraction of $v \in N(u)$ we have that $\mathbf{Inf}_{\pi_{uv}(i)}(T_{1-\delta} f_v) \geqslant \varepsilon/2$. Let

$$L_v = \{j \in [L] \mid \mathbf{Inf}_j(T_{1-\delta} f_v) \geqslant \varepsilon/2\}.$$

We know that $|L_v| \leqslant \frac{2}{\varepsilon \cdot e \cdot \ln(1-\delta)}$ using Lemma 2.5. Consider the following randomized labeling to the Unique Games instance. For each $u \in U$, if there exists $i \in [L]$ such that $\mathbf{Inf}_i(T_{1-\delta} f_u) \geqslant \varepsilon$ then assign label i to u. Otherwise, assign a random label from $[L]$ to u. For each $v \in V$, pick a random label from L_v if it is non-empty. If $|L_v| = 0$ then pick a random label from $[L]$. The randomized labeling satisfies at least $\eta \cdot \frac{\varepsilon}{2} \cdot \frac{1}{|L_v|} \geqslant \eta \cdot \frac{\varepsilon}{2} \cdot \frac{\varepsilon \cdot e \cdot \ln(1-\delta)}{2} = \eta'$ fraction of the edges in expectation, which is a contradiction. □

Let $U' \subseteq U$ be the set of $u \in U$ for which the above claim holds. Using the above claim, we have

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_1}(\mathcal{S}) &= 2 \cdot \nu_{q_\star}(f) - 2 \mathop{\mathbf{E}}_{u \in U} \mathop{\mathbf{E}}_{(x,y) \sim \mathcal{A}_{\rho_1, q_\star}^{\otimes L}} [f_u(x) f_u(y)] \\
&\leqslant 2 \cdot \nu_{q_\star}(f) - 2 \left((1-\eta) \mathop{\mathbf{E}}_{u \in U'} \left[\overline{\Gamma}_{\rho_1}(\nu_{q_\star}^u(f), \nu_{q_\star}^u(f)) - \frac{\eta}{2} \right] + \eta \cdot 0 \right) \\
&\leqslant 2 \cdot \nu_{q_\star}(f) - 2 \mathop{\mathbf{E}}_{u \in U'} [\overline{\Gamma}_{\rho_1}(\nu_{q_\star}^u(f), \nu_{q_\star}^u(f))] + \eta.
\end{aligned}$$

Now using the convexity of the function $\overline{\Gamma}_\rho(x, y)$, we have

$$\begin{aligned}
\mathop{\mathbf{E}}_{u \in U'} \left[\overline{\Gamma}_{\rho_1}(\nu_{q_\star}^u(f), \nu_{q_\star}^u(f)) \right] &\geqslant \overline{\Gamma}_{\rho_1} \left(\mathop{\mathbf{E}}_{u \in U'}(\nu_{q_\star}^u(f)), \mathop{\mathbf{E}}_{u \in U'}(\nu_{q_\star}^u(f)) \right) \\
&\geqslant \overline{\Gamma}_{\rho_1} \left(\nu_{q_\star}(f) - \eta, \nu_{q_\star}(f) - \eta \right)
\end{aligned}$$

where the last inequality follows from $\left| \mathbf{E}_{u \in U}[\nu_{q_\star}^u(f)] - \mathbf{E}_{u \in U'}[\nu_{q_\star}^u(f)] \right| \leqslant \eta$ and the fact that $\overline{\Gamma}_\rho(x, y)$ is an increasing function of x and y. Thus, we have

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_1}(\mathcal{S}) &\leqslant 2 \cdot \nu_{q_\star}(f) - 2 \cdot \overline{\Gamma}_{\rho_1} \left(\nu_{q_\star}(f) - \eta, \nu_{q_\star}(f) - \eta \right) + \eta \\
&\leqslant 2 \cdot \nu_{q_\star}(f) - 2 \cdot \overline{\Gamma}_{\rho_1} \left(\nu_{q_\star}(f), \nu_{q_\star}(f) \right) + 3\eta. \qquad (2)
\end{aligned}$$

The exact same calculation shows that

$$\mathbf{Cut}_{\mathcal{G}_2}(\mathcal{S}) \leqslant 2 \cdot \nu_{(1-q_\star)}(f) - 2 \cdot \overline{\Gamma}_{\rho_1} \left(\nu_{(1-q_\star)}(f), \nu_{(1-q_\star)}(f) \right) + 3\eta. \qquad (3)$$

We now analyze the value of the cut given by f in $\mathcal{G}_3$:

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_3}(\mathcal{S}) &= \mathop{\mathbf{E}}_{u \in U} \mathop{\mathbf{E}}_{(x,y) \sim \mathcal{B}_{\rho_2, q_\star}^{\otimes L}} [f_u(x) + f_u(y) - 2 f_u(x) f_u(y)] \\
&= \nu_{q_\star}(f) + \nu_{(1-q_\star)}(f) - 2 \mathop{\mathbf{E}}_{u \in U} \mathop{\mathbf{E}}_{(x,y) \sim \mathcal{B}_{\rho_2, q_\star}^{\otimes L}} [f_u(x) f_u(y)].
\end{aligned}$$

Similar to Claim 4.3, we have,

Claim 4.4. *For at least $(1-\eta)$ fraction of $u \in U$,*

$$\mathop{\mathbf{E}}_{(x,y)\sim\mathcal{B}_{\rho_1,q_\star}^{\otimes L}}[f_u(x)f_u(y)] \geqslant \overline{\Gamma}_{\rho_2}\left(\nu_{q_\star}^u(f), \nu_{(1-q_\star)}^u(f)\right) - \frac{\eta}{2}.$$

Proof. The proof is similar to the proof of Claim 4.3 once we conclude, using Corollary 2.11 that there exists $\varepsilon, \delta > 0$ such that for at least η fraction of f_u we have that $\mathbf{Inf}_i(T_{1-\delta}f_u) \geqslant \varepsilon$ for some $i \in [L]$. □

Let $U'' \subseteq U$ be the set of $u \in U$ for which the above claim holds. Using the above claim, we have

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_3}(\mathcal{S}) &= \nu_{q_\star}(f) + \nu_{(1-q_\star)}(f) - 2\mathop{\mathbf{E}}_{u\in U}\mathop{\mathbf{E}}_{(x,y)\sim\mathcal{B}_{\rho_2,q_\star}^{\otimes L}}[f_u(x)f_u(y)] \\
&\leqslant \nu_{q_\star}(f) + \nu_{(1-q_\star)}(f) - 2\left((1-\eta)\mathop{\mathbf{E}}_{u\in U''}\left[\overline{\Gamma}_{\rho_2}\left(\nu_{q_\star}^u(f), \nu_{(1-q_\star)}^u(f)\right) - \frac{\eta}{2}\right] + \eta\cdot 0\right) \\
&\leqslant \nu_{q_\star}(f) + \nu_{(1-q_\star)}(f) - 2\mathop{\mathbf{E}}_{u\in U}\left[\overline{\Gamma}_{\rho_2}\left(\nu_{q_\star}^u(f), \nu_{(1-q_\star)}^u(f)\right)\right] + \eta.
\end{aligned}$$

Again, using the convexity of $\overline{\Gamma}_{\rho_2}$,

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_3}(\mathcal{S}) &\leqslant \nu_{q_\star}(f) + \nu_{(1-q_\star)}(f) - 2\overline{\Gamma}_{\rho_2}\left(\nu_{q_\star}(f) - \eta, \nu_{(1-q_\star)}(f) - \eta\right) + \eta \\
&\leqslant \nu_{q_\star}(f) + \nu_{(1-q_\star)}(f) - 2\overline{\Gamma}_{\rho_2}\left(\nu_{q_\star}(f), \nu_{(1-q_\star)}(f)\right) + 3\eta \qquad (4)
\end{aligned}$$

Now, let us compare the solution w.r.t $2(1-q_\star) - \eta$. For the notational convenience let $\nu_1 = \nu_{q_\star}(f)$ and $\nu_2 = \nu_{(1-q_\star)}(f)$. Then,

$$\begin{aligned}
\mathbf{Cut}_{\mathcal{G}_1}(\mathcal{S}) &\leqslant 2\cdot\nu_1 - 2\overline{\Gamma}_{\rho_1}(\nu_1,\nu_1) + 3\eta \\
\mathbf{Cut}_{\mathcal{G}_2}(\mathcal{S}) &\leqslant 2\cdot\nu_2 - 2\overline{\Gamma}_{\rho_1}(\nu_2,\nu_2) + 3\eta \\
\mathbf{Cut}_{\mathcal{G}_3}(\mathcal{S}) &\leqslant \nu_1 + \nu_2 - 2\overline{\Gamma}_{\rho_2}(\nu_1,\nu_2) + 3\eta.
\end{aligned}$$

In this case, ν_1, ν_2 are the free parameters which come from the indicator function f of the cut we started with. Define the following ranges:

$$\begin{aligned}
R_1(\eta) &= \mathop{\texttt{range}}_{\nu_1\in[0,1]}\left\{\frac{2\nu_1 - 2\Gamma_{\rho_1}(\nu_1,\nu_1) + 3\eta}{2(1-q_\star) - \eta} \geqslant (\alpha_{GW} - 10^{-5})\right\}, \\
R_2(\eta) &= \mathop{\texttt{range}}_{\nu_2\in[0,1]}\left\{\frac{2\nu_2 - 2\Gamma_{\rho_1}(\nu_2,\nu_2) + 3\eta}{2(1-q_\star) - \eta} \geqslant (\alpha_{GW} - 10^{-5})\right\}, \\
R_3(\eta) &= \mathop{\texttt{range}}_{\nu_1,\nu_2\in[0,1]}\left\{\frac{\nu_1 + \nu_2 - 2\overline{\Gamma}_{\rho_2}(\nu_1,\nu_2) + 3\eta}{2(1-q_\star) - \eta} \geqslant (\alpha_{GW} - 10^{-5})\right\}.
\end{aligned}$$

If we want to get a cut with values $(\alpha_{GW} - 10^{-5}) \cdot (2(1-q_\star) - \eta)$ in all the graphs $\mathcal{G}_1, \mathcal{G}_2$ and $\mathcal{G}_3$ then we must have the $(R_1(\eta) \times R_2(\eta)) \cap R_3(\eta) \neq \emptyset$.

By performing numerical calculations, we show that there exists an absolute constant $\eta_0 > 0$ such that for all $0 < \eta \leqslant \eta_0$, $(R_1(\eta) \times R_2(\eta)) \cap R_3(\eta)$ is in fact $\emptyset$. Thus, no matter which densities $\nu_1 = \nu_{q_\star}(f)$ and $\nu_2 = \nu_{(1-q_\star)}(f)$ we choose, there exists an $i \in [3]$ such that the value of the cut in graph $\mathcal{G}_i$ given by f will be less than $(\alpha_{GW} - \varepsilon_0)(2(1-q_\star) - \eta)$ for some fixed constant $\varepsilon_0 \geqslant 10^{-5}$. $\square$

Acknowledgement : Our numerical calculations involve minor modifications of the prover code[1] written by Austrin et al. [ABG16] which uses interval arithmetic to get a computer generated proof. We are indebted to the authors of [ABG16] for making it available online.

References

[ABG16] Per Austrin, Siavosh Benabbas, and Konstantinos Georgiou. Better balance by being biased: A 0.8776-approximation for max bisection. *ACM Transactions on Algorithms (TALG)*, 13(1):2, 2016.

[AKS11] Per Austrin, Subhash Khot, and Muli Safra. Inapproximability of Vertex Cover and Independent Set in Bounded Degree Graphs. *Theory of Computing*, 7(3):27–43, 2011.

[AS19] Per Austrin and Aleksa Stankovic. Global Cardinality Constraints Make Approximating Some Max-2-CSPs Harder. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2019, September 20-22, 2019, Massachusetts Institute of Technology, Cambridge, MA, USA.*, pages 24:1–24:17, 2019.

[BKK+18] Amey Bhangale, Subhash Khot, Swastik Kopparty, Sushant Sachdeva, and Devanathan Thiruvenkatachari. Near-optimal approximation algorithm for simultaneous Max-Cut. In *Proceedings of the Twenty-Ninth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2018, New Orleans, LA, USA, January 7-10, 2018*, pages 1407–1425, 2018.

[BKS15] Amey Bhangale, Swastik Kopparty, and Sushant Sachdeva. Simultaneous Approximation of Constraint Satisfaction Problems. In *Automata, Languages, and Programming - 42nd International Colloquium, ICALP 2015, Kyoto, Japan, July 6-10, 2015, Proceedings, Part I*, pages 193–205, 2015.

[1] available at `https://github.com/austrin/max-bisection-analysis`

[DMR09] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *SIAM Journal on Computing*, 39(3):843–873, 2009.

[GW95] Michel X Goemans and David P Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *Journal of the ACM (JACM)*, 42(6):1115–1145, 1995.

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the thiry-fourth annual ACM symposium on Theory of computing*, pages 767–775. ACM, 2002.

[KKMO07] 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, April 2007.

[MOO05] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 21–30. IEEE, 2005.

[Mos10] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19(6):1713–1756, 2010.

ECCC ISSN 1433-8092

https://eccc.weizmann.ac.il