Hardness of Approximating Bounded-Degree Max 2-CSP and Independent Set on k-Claw-Free Graphs

Euiwoong Lee
University of Michigan, Ann Arbor, MI, USA

Pasin Manurangsi
Google Research, Bangkok, Thailand

Abstract

We consider the question of approximating Max 2-CSP where each variable appears in at most d constraints (but with possibly arbitrarily large alphabet). There is a simple $(\frac{d+1}{2})$-approximation algorithm for the problem. We prove the following results for any sufficiently large d:

- Assuming the Unique Games Conjecture (UGC), it is NP-hard (under randomized reduction) to approximate this problem to within a factor of $\left(\frac{d}{2} - o(d)\right)$.
- It is NP-hard (under randomized reduction) to approximate the problem to within a factor of $\left(\frac{d}{3} - o(d)\right)$.

Thanks to a known connection [15], we establish the following hardness results for approximating Maximum Independent Set on k-claw-free graphs:

- Assuming the Unique Games Conjecture (UGC), it is NP-hard (under randomized reduction) to approximate this problem to within a factor of $\left(\frac{k}{4} - o(k)\right)$.
- It is NP-hard (under randomized reduction) to approximate the problem to within a factor of $\left(\frac{k}{3+2\sqrt{2}} - o(k)\right) \geq \left(\frac{k}{5.829} - o(k)\right)$.

In comparison, known approximation algorithms achieve $\left(\frac{k}{2} - o(k)\right)$-approximation in polynomial time [32, 37] and $(\frac{k}{3} + o(k))$-approximation in quasi-polynomial time [11].

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness; Theory of computation → Approximation algorithms analysis

Keywords and phrases Hardness of Approximation, Bounded Degree, Constraint Satisfaction Problems, Independent Set

Digital Object Identifier 10.4230/LIPIcs.ITCS.2024.71

Funding *Euiwoong Lee*: Supported in part by NSF grant CCF-2236669 and Google.

Acknowledgements This work was initiated at Dagstuhl Seminar 23291 "Parameterized Approximation: Algorithms and Hardness". We thank the organizers and participants of the workshop for helpful discussions.

1 Introduction

An instance of Max 2-CSP consists of variables–each of which can take a value of an alphabet–together with constraints, each involving a pair of variables. The goal is to find an assignment to the variables that satisfies as many constraints as possible. The Max 2-CSP problem is a cornerstone of the field of hardness of approximation as it[1] is often used as a starting point in hardness of approximation reductions. When the constraints are restricted to certain predicates–such as 3SAT or Max-Cut, tight hardness of approximation results are known through a series of influential work (e.g. [22, 25]). In fact, it is known that a

[1] Or more precisely, its special case known as *Label Cover* or *Projection Games*

certain semi-definite program relaxation provides essentially the best approximation ratio achievable in polynomial time [34]. Meanwhile, besides the predicate, there are also other parameters that can affect the approximation ratio. One of which is the *(maximum) degree* of the instance d, defined as the maximum number of constraints that a variable appears in. A number of previous studies have investigated how the degree affects the approximation ratio (e.g. [21, 38, 27, 3]), partly because, as we will see in more detail below, it affects the (in)approximation ratio of subsequent problems in hardness reductions. In this work, we focus on determining the approximation ratio in terms of this parameter d alone (regardless of the predicate or other parameters) and ask:

What is the best possible approximation ratio for Max 2-CSP in terms of d?

Regarding this question, there is a simple $(\frac{d+1}{2})$-approximation algorithm for the problem (see Appendix B). On the hardness front, Laekhanukit [27] proved NP-hardness (under randomized reduction) with a factor of $\Omega(d/\log d)$ for any sufficiently large d. Furthermore, under a less standard "Strongish Planted Clique Hypothesis"[2], a hardness of a factor $\Omega(d)$ is known but without any explicit constant in the inapproximability factor[3] [30].

Maximum Independent Set on k-Claw-Free Graphs

While the Maximum Independent Set problem is well known to be NP-hard to approximate to within a factor of $n^{1-\epsilon}$ where n is the number of vertices [20, 39], there are multiple classes of graphs for which this can be significantly improved upon. One such class is that of *k-claw-free* graphs. Recall that a k-claw (i.e. $K_{1,k}$) is the star graph with a center vertex connecting to k other vertices (where there are no edges between these k vertices). A graph is k-claw-free if it does not contain a k-claw as an *induced* subgraph. The classic local search algorithm of Berman [4] achieves $(\frac{k}{2}+\epsilon)$-approximation in polynomial time for any constant $\epsilon > 0$. Recently, this has been improved by [32, 37] to achieve a slightly-better-than-$(k/2)$ approximation ratio; in particular, [37] achieves approximation ratio of $(\frac{k}{2}-\frac{1}{3}+o(1))$ where $o(1)$ is a term that converges to 0 as $k \to \infty$. Meanwhile, in quasi-polynomial (i.e. $n^{O(\log n)}$) time, this ratio can be improved[4] to $\left(\frac{k}{3}+\epsilon\right)$ for any constant $\epsilon > 0$ [11]. These algorithms are based on local search approaches. Meanwhile, several works have also investigated the power of LP/SDP relaxations of the problem: Chudnovsky and Seymour [8] showed that a standard SDP relaxation yields 2-approximation when $k = 3$, but a recent work by Chalermsook et al. [6] shows large integrality gaps for $k > 3$.

On the hardness of approximation front, Hazan et al. [23] proved that the problem is NP-hard to approximate to within $\Omega(k/\log k)$ factor. In a recent work, Dvorak et al. [15] observed that the classic FGLSS reduction [16] provides an approximation-preserving reduction from Max 2-CSP with maximum degree d to maximum independent set in $2d$-claw-free graphs. This reduction, together with the aforementioned hardness from [27], gives an alternative NP-hardness proof with a similar inapproximability factor. Meanwhile, plugging this to the

---

[2] The Strongish Planted Clique Hypothesis states that no $n^{o(\log n)}$-time algorithm can distinguish between a $G(n, 1/2)$ random graph and one in which a clique of size n^c is planted for some absolute constant $c > 0$.

[3] In fact, this hardness only states that, for each $c_1 > 0$, there exists $c_2 > 0$ such that no $O(n^{c_1})$-time algorithm achieves $c_2 d$-approximation ratio. In other words, it does not rule out e.g. $n^{O(\log^* n)}$ time algorithm from achieving $o(d)$-approximation ratio.

[4] We only discuss the unweighted case in our paper as our hardness results apply to this case; for the weight case, it is not known how to achieve $\left(\frac{k}{3}+o(k)\right)$-approximation in quasi-polynomial time. Please refer to [33] for the best approximation algorithms known for the weighted case.

other aforesaid result of [30] implies that no polynomial-time algorithm can achieves $o(k)$ approximation ratio (albeit without an explicit constant again) under the Strongish Planted Clique Hypothesis. [15] also consider parameterization based on the independent set size and prove several hardness results in that setting; we defer the discussion on this to Appendix C.

1.1 Our Results

Bounded-Degree Max 2-CSP

Our main contribution is a nearly tight hardness of approximation result for Max 2-CSP in terms of d assuming the Unique Games Conjecture (UGC) [24]:

▶ **Theorem 1.** *Assuming the UGC, for any $\epsilon \in (0, 1/2)$, there exists $d_0 \in \mathbb{N}$ such that the following holds for every positive integer $d \geq d_0$: Unless NP = BPP, there is no polynomial-time $d(1/2 - \epsilon)$-approximation algorithm for d-bounded-degree Max 2-CSP.*

As stated earlier, there is a simple $(\frac{d+1}{2})$-approximation algorithm and thus our result is within a factor of $1 + o(1)$ of this upper bound (as $d \to \infty$). To the best of our knowledge, this is also the first $\Omega(d)$ hardness of approximation result with an explicit constant for the problem (under any assumption). For NP-hardness, we prove a slightly weaker result where the factor is instead $\approx d/3$:

▶ **Theorem 2.** *For any $\epsilon \in (0, 1/3)$, there exists $d_0 \in \mathbb{N}$ such that the following holds for every positive integer $d \geq d_0$: Unless NP = BPP, there is no polynomial-time $d(1/3 - \epsilon)$-approximation algorithm for d-bounded-degree Max 2-CSP.*

Independent Set in Claw-Free Graphs

Leveraging the connection between bounded-degree Max 2-CSP and Maximum Independent Set in claw-free graphs [15] discussed above, we arrive at a $\approx k/4$ hardness for the latter, assuming the Unique Games Conjecture.

▶ **Theorem 3.** *Assuming the UGC, for any $\epsilon \in (0, 1/4)$, there exists $k_0 \in \mathbb{N}$ such that the following holds for every positive integer $k \geq k_0$: Unless NP = BPP, there is no polynomial-time $k(1/4 - \epsilon)$-approximation algorithm for Maximum Independent Set on k-claw-free graphs.*

Again, this is the first $\Omega(k)$ hardness for the problem with an explicit constant. Furthermore, this is within a factor of 2 (as $k \to \infty$) of the aforementioned polynomial-time approximation algorithms [32, 37] and within a factor of $4/3 + o(1)$ of the quasi-polynomial time approximation algorithm [11].

For NP-hardness result, we get a slightly weaker factor that is $\approx k/5.829$ instead:

▶ **Theorem 4.** *For any $\epsilon \in \left(0, \frac{1}{3+2\sqrt{2}}\right)$, there exists $k_0 \in \mathbb{N}$ such that the following holds for every positive integer $k \geq k_0$: Unless NP = BPP, there is no polynomial-time $k\left(\frac{1}{3+2\sqrt{2}} - \epsilon\right)$-approximation algorithm for Maximum Independent Set on k-claw-free graphs.*

1.2 Technical Overview

We now briefly (and informally) discuss our techniques. Perhaps surprisingly, we use the same strategy as in previous work: sparsify a dense(er) 2-CSP instance by randomly sampling its constraints. This strategy–originated in [38]–has been used in many subsequent papers on the topic (e.g. [18, 27, 12, 28]). As we will elaborate more below, the main "twist" in our work is that, instead of starting with Max 2-CSP hardness with a gap roughly similar to the desired gap after sampling, we start with Max 2-CSP hardness with a much larger gap.

To discuss this in more detail, let us first recall the standard subsampling procedure. We start with a 2-CSP instance Π and produces Π' as follows: (i) keep each edge in Π with probability p and (ii) remove edges until the every vertex has degree at most d. For simplicity of presentation, let us assume that the constraint graph in Π is $\tilde{d}$-regular. In this case, we can let $p = d/\tilde{d}$. The completeness of the reduction is obvious: if Π is fully satisfiable[5], then Π' is also fully satisfiable.

The main challenge is in analyzing the soundness. Again, suppose for simplicity that we did not have to apply step (ii). Suppose that any assignment satisfies less than γ fraction of constraints in Π. What can we say about Π'?

A standard soundness argument here is to use a concentration bound to show that, for some $\gamma' > \gamma$ and any fixed assignment ψ, the probability that ψ satisfies more than γ' fraction of constraints in Π' is at most q. Then, using a union bound over all assignments, one arrive at a conclusion that no assignment satisfies more than γ' fraction of constraints in Π'. This gives a gap of γ'. Recall that we want $\gamma' = \Omega(1)/d$. Note also that there are R^n assignments, where R denote the alphabet size and n denote the number of variables. Therefore, we need $q \ll R^{-n}$ for this argument to work. Meanwhile, when $\gamma = (\gamma')^{\omega(1)}$ the multiplicative Chernoff bound gives

$$q \leq O\left(\frac{\gamma'}{\gamma}\right)^{\gamma' dn/2} = (1/\gamma)^{-\gamma' dn/2 \cdot (1-o(1))}.$$

Comparing this with the required $q \ll R^{-n}$, it suffices for us to take $\gamma' = \frac{2}{d} \cdot \log_{1/\gamma} R \cdot (1+o(1))$. Putting it differently, if we start with a hardness for Max 2-CSP with a gap of $1/\gamma = R^{\nu}$, then we end up with a gap of $\frac{d}{2} \cdot \nu \cdot (1-o(1))$. Under the UGC, we show that such a hardness can be proved for $\nu = 1 - o(1)$. (See Appendix A.) This immediately yields Theorem 1.

It is crucial to point out that we require $\gamma \ll \gamma'$ as otherwise, if $\gamma' = \Theta(\gamma)$, the bound would only be $\exp\left(-\Omega\left(\gamma' dn\right)\right)$ which would require us to take $\gamma' = \Omega\left(\frac{\log R}{d}\right)$. This $\log R$ factor is essentially what differentiates us from previous work on similar topics, such as [27].

Optimizing Parameters for NP-hardness

For NP-hardness, we start with the NP-hardness result of [7] where $\gamma = R^{1/2-o(1)}$ or $\nu = 1/2 - o(1)$. If we were to plug this into the above argument directly, we would get a gap of only $\frac{d}{4} \cdot (1 - o(1))$. We are able to get a better gap of $\frac{d}{3} \cdot (1 - o(1))$ by observing that the instance of [7] is bipartite and has RHS alphabet of size only $\sqrt{R}$. This allows us to use a union bound on only $R^{3n/4}$ assignments (instead of R^n), which improves the inapproximability ratio as claimed.

Independent Set on k-Claw-Free Graphs

For UGC-hardness of of Maximum Independent Set on k-claw-free graphs, we can combine the UGC-hardness for Max 2-CSP with bounded degree (Theorem 1) together with the aforementioned connection from [15], which immediately yields Theorem 3. As for NP-hardness, the same strategy only gives us $\frac{k}{6} \cdot (1 - o(1))$ inapproximability factor. To improve

[5] Again, this is for simplicity; in the actual reduction, we only have almost satisfiability.

this, we observe that the observation in [15] can be further refined when the graph is bipartite and the maximum degree on each sides are different (Lemma 12). By balancing these degree parameters (namely letting the LHS degree being $\approx \sqrt{2}$ times that of the RHS), we arrive at the claimed $\frac{k}{3+2\sqrt{2}} \cdot (1 - o(1))$ hardness factor.

1.3 Other Related Work

Maximum Independent Set on k-claw-free graphs is closely related to many other important problems in literature. For example, it contains Maximum Independent Set on bounded-degree graphs (where the maximum degree is at most k) as a special case. It turns out that the latter is easier: a $\tilde{O}(k/\log^2 k)$-approximation algorithm [2] is known and this is essentially tight [1]. Another closely related problem is the k-Set Packing problem, in which we are given sets of size at most k and would like to pick as many disjoint sets as possible. It is simple to see that the problem is equivalent to finding an independent set in the "conflict graph"–where each set becomes a vertex and two vertices are linked if and only if the sets intersect–and that this conflict graph is $(k+1)$-claw-free. Thus, all aforementioned approximation algorithms for Maximum Independent Set on claw-free graphs immediately apply to k-Set Packing. However, the latter can also be less challenging: the aforementioned quasi-polynomial time algorithm of Cygan et al. [11] for the former can be sped up to run in polynomial time while acheiving a similar approximation ratio [9, 36]. Meanwhile, the best known hardness of approximation for the problem remains an NP-hardness with approximation factor $\Omega(k/\log k)$ due to Hazan et al. [23].

2 Preliminaries

We use $\mathrm{indep}(G)$ to denote the size of the maximum independent set in the graph G.

2.1 Concentration Inequalities

We recall the standard multiplicative Chernoff bound:

▶ **Theorem 5** (Multiplicative Chernoff Bound). *Let $X_1, \dots, X_m$ be i.i.d. Bernoulli random variable with mean at most μ, and let $S = X_1 + \cdots + X_m$. Then, for any $\theta > \mu m$, we have*

$$\Pr\left[S > \theta\right] < \exp(\theta - \mu m) \left(\frac{\mu m}{\theta}\right)^{\theta}$$

For $\tau > 0$ and $x \in \mathbb{R}$, let $\mathrm{clip}_\tau(x) := \min\{x, \tau\}$. We will need the following lemma for the purpose of analyzing edge removal to bound the maximum degree of a random subgraph.

▶ **Lemma 6.** *Let $X_1, \dots, X_m$ be i.i.d. Bernoulli random variable with mean at most μ, and let $S = X_1 + \cdots + X_m$. Then, for any integer $\tau > \mu m$, we have*

$$\mathbb{E}\left[S - \mathrm{clip}_\tau(S)\right] \leq \left(\frac{\mu m}{\tau - \mu m}\right)^2.$$

Proof. For any $i \in [m]$, we have $\mathrm{Var}[X_1 + \cdots + X_{i-1}] = \mathrm{Var}[X_1] + \cdots \mathrm{Var}[X_{i-1}] \leq \mu(i-1)$.
We can rewrite the LHS as

$$\mathbb{E}[S - \mathrm{clip}_\tau(S)] = \mathbb{E}\left[\sum_{i \in [m]} X_i \cdot \mathbf{1}[X_1 + \cdots X_{i-1} \geq \tau]\right]$$

$$
\begin{aligned}
&= \sum_{i \in [m]} \mathbb{E}[X_i] \cdot \Pr[X_1 + \cdots X_{i-1} \geq \tau] \\
&\leq \sum_{i \in [m]} \mu \cdot \frac{\mu(i-1)}{(\tau - \mu(i-1))^2} \\
&\leq \left(\frac{\mu m}{\tau - \mu m}\right)^2,
\end{aligned}
$$

where in the first inequality we use Chebyshev's inequality. ◀

2.2 Constraint Satisfaction Problems

Formal definitions of a 2-CSP instance and its assignment are given below.

▶ **Definition 7.** *A 2-CSP instance* Π *consists of:*

- Constraint graph $G = (V, E)$.
- Alphabet Σ_v *for all* $v \in V$.
- *For each* $e = (u, v) \in E$, *a* constraint $R_e \subseteq \Sigma_u \times \Sigma_v$.

An assignment *is a tuple* $(\psi_v)_{v \in V}$ *such that* $\psi_v \in \Sigma_v$. *Its value* $\mathrm{val}_\Pi(\psi)$ *is defined as the fraction of edges* $e = (u, v) \in E$ *such that* $(\psi_u, \psi_v) \in E$*; such an edge (or constraint) is said to be satisfied. The value of the instance is defined as* $\mathrm{val}(\Pi) = \max_\psi \mathrm{val}_\Pi(\psi)$ *where the maximum is over all assignments* ψ.

Additionally, we use the following terminologies for CSPs:

- A 2-CSP instance is *d-bounded-degree* if the every vertex in the constraint graph G has degree at most d.
- The *alphabet size* of a 2-CSP instance is $\max_{v \in V} |\Sigma_v|$.
- A 2-CSP instance is *bipartite* if the constraint graph $G = (A, B, E)$ is a bipartite graph.
- A bipartite 2-CSP instance is (d_1, d_2)*-biregular* if every left-hand side vertex (in A) has degree d_1 and every right-hand side vertex (in B) has degree d_2.
- A bipartite 2-CSP instance is (d_1, d_2)*-bounded-degree* if every left-hand side vertex (in A) has degree at most d_1 and every right-hand side vertex (in B) has degree at most d_2.
- The left (resp. right) alphabet size of a bipartite 2-CSP instance is $\max_{a \in A} |\Sigma_a|$ (resp. $\max_{b \in B} |\Sigma_b|$).

2.3 Hardness of 2-CSP in terms of Alphabet Size

As discussed in the introduction, we need hardness of almost-perfect completeness 2-CSP with a gap that is polynomial in the alphabet size. For NP-hardness, the best known result is due to [7], which has a gap of $R^{1/2-o(1)}$:

▶ **Theorem 8** ([7])**.** *For any* $\zeta > 0$ *and sufficiently large* $R \in \mathbb{N}$ *such that* $\sqrt{R}$ *is a prime number, there exists* $d_1, d_2 \in \mathbb{N}$ *such that it is NP-hard, given a bipartite* (d_1, d_2)*-biregular*[6] *2-CSP* Π *with left alphabet size* R *and right alphabet size* $\sqrt{R}$, *to distinguish between the following two cases:*

- *(Yes Case)* $\mathrm{val}(\Pi) \geq 1 - \zeta$,
- *(No Case)* $\mathrm{val}(\Pi) \leq O\left(\frac{\log R}{\sqrt{R}}\right)$.

[6] Note that biregularity follows immediately if we intiate the reduction of [7] with a biregular Label Cover.

For UGC-hardness, a standard proof technique by Khot, Kindler, Mossel, and O'Donnell [25] yields a hardness of factor $R^{1-o(1)}$, as stated below. Since we are not aware[7] of such a result fully written down in literature, we provide its proof in Appendix A for completeness.

▶ **Theorem 9.** *Assuming the Unique Games Conjecture, for any $\zeta > 0$ and sufficiently large $R \in \mathbb{N}$, there exists $d_1, d_2 \in \mathbb{N}$ such that it is NP-hard, given a bipartite (d_1, d_2)-biregular 2-CSP Π with alphabet size R, to distinguish between the following two cases:*
- *(Yes Case)* $\text{val}(\Pi) \geq 1 - \zeta$,
- *(No Case)* $\text{val}(\Pi) \leq O\left(\frac{\log^2 R}{R}\right)$.

3 Hardness of Bounded-Degree 2-CSP

In this section, we present our main reduction and prove Theorem 1 and Theorem 2.

3.1 Adjusting the Degrees

As alluded to in the introduction, it will be useful to have a flexible control of the degrees of the two sides of the constraint graph. This can be easily done by copying the vertices on each side, as formalized below.

▶ **Lemma 10.** *For any $d_1, d_2, c_1, c_2 \in \mathbb{N}$, there is a polynomial-time reduction from a bipartite (d_1, d_2)-biregular 2-CSP Π to a bipartite $(c_2 d_1 d_2, c_1 d_1 d_2)$-biregular 2-CSP Π' such that $\text{val}(\Pi') = \text{val}(\Pi)$. Moreover, the reduction preserves the left and right alphabet sizes.*

Proof. Let the original 2-CSP instance be $\Pi = (G = (A, B, E), (\Sigma_v)_{v \in A \cup B}, (R_e)_{e \in E})$ where G is (d_1, d_2)-biregular. We define $\Pi = (G' = (A', B', E'), (\Sigma_{v'})_{v' \in A' \cup B'}, (R_{e'})_{e' \in E'})$ where
- $A' = A \times [d_1] \times [c_1]$,
- $B' = B \times [d_2] \times [c_2]$,
- $E' = \{((a, i_1, j_1), (b, i_2, j_2)) \mid (a, b) \in E, i_1 \in [d_1], i_2 \in [d_2], j_1 \in [c_1], j_2 \in [c_2]\}$,
- $\Sigma'_{(v,i,j)} = \Sigma_v$ for all $(v, i, j) \in A' \cup B'$.
- $R_{((a,i_1,j_1),(b,i_2,j_2))} = R_{(a,b)}$ for all $((a, i_1, j_1), (b, i_2, j_2)) \in E'$.

To see that $\text{val}(\Pi') \geq \text{val}(\Pi)$, let ψ denote the assignment of Π with $\text{val}_\Pi(\psi) = \text{val}(\Pi)$. Define an assignment ψ' of Π' such that $\psi'_{(v,i,j)} := \psi_v$. it is simple to check that $\text{val}_{\Pi'}(\psi') = \text{val}_\Pi(\psi) = \text{val}(\Pi)$.

On the other hand, to see that $\text{val}(\Pi') \leq \text{val}(\Pi)$, notice that E' is can be partitioned into $E_{(i_1,j_1,i_2,j_2)} := \{((a, i_1, j_1), (b, i_2, j_2)) \mid (a, b) \in E\}$ where $i_1 \in [d_1], i_2 \in [d_2], j_1 \in [c_1], j_2 \in [c_2]$. Thus, since any assignment satisfies at most $\text{val}(\Pi)$ fraction of $E_{(i_1,j_1,i_2,j_2)}$, we can conclude that any assignment also satisfies at most $\text{val}(\Pi)$ fraction of E. ◀

3.2 Main Reduction: Degree Reduction via Subsampling

We are now ready to state our main reduction and its properties. For readers interested in only the UGC-hardness results, it suffices to think of just the case where the degree bounds d_A, d_B are equal, the alphabet sizes are equal (i.e. $t = 1$) and $\nu = 1 - o(1)$ in the theorem statement below.

[7] While Kindler et al. [26] proved a UGC-hardness result with a similar factor, their result does not satisfy almost-perfect completeness, making it unsuitable for our purpose.

► **Theorem 11.** *For any $t, \delta, \nu \in (0,1]$ such that $\delta < \nu$, any positive integer C, and any sufficiently large positive integers $d_A, d_B \geq d_0(\delta, \nu)$ and $R \geq R_0(\delta, \nu, t, d_A, d_B)$, the following holds: there is a randomized polynomial-time reduction from a bipartite $(d_A C, d_B C)$-biregular 2-CSP Π' with left alphabet size at most R and right alphabet size at most R^t to a (d_A, d_B)-bounded-degree 2-CSP Π'' such that, with probability $2/3$, we have*
- *(Completeness)* $\mathrm{val}(\Pi'') \geq \mathrm{val}(\Pi') - \delta$*, and,*
- *(Soundness) If* $\mathrm{val}(\Pi') \leq \frac{1}{R^\nu}$*, then* $\mathrm{val}(\Pi'') \leq \frac{1}{\nu-\delta}\left(\frac{1}{d_A} + \frac{t}{d_B}\right)$.

Proof. Let $\Pi' = (G' = (A', B', E'), (\Sigma_{v'})_{v' \in A' \cup B'}, (R_e)_{e \in E'})$ denote the original instance. We select the parameters as follows:
- $\lambda := 0.001 \min\{\delta, \nu\}$,
- $p := \frac{1-\lambda}{C}$,
- $d_0 = \frac{10000}{\lambda^3}$,
- $\chi := \frac{1}{\nu - 2\lambda}\left(\frac{1}{d_A} + \frac{t}{d_B}\right)$,
- $R_0 = \max\left\{\left(\frac{e}{\chi}\right)^{1/\lambda}, 100^{1/\left(\frac{1}{d_A} + \frac{t}{d_B} - (\nu-\lambda)\cdot\chi\right)}\right\}$.
- $n_E := d_A|A'|$.

We construct the instance Π'' as follows:
1. First, let $E_1 \subseteq E'$ be a subset of edges where each edge in E' is kept with probability p.
2. Let $E'' = E_1$ and $G'' = (A', B', E'')$.
3. For all $a \in A'$: If $\deg_{G''}(a) > d_A$, remove (arbitrary) $d_A - \deg_{G''}(a)$ edges adjacent to a from E''.
4. For all $b \in B'$: If $\deg_{G''}(b) > d_B$, remove (arbitrary) $d_B - \deg_{G''}(b)$ edges adjacent to b from E''.
5. Let Π'' be $(G'' = (A', B', E''), (\Sigma_{v'})_{v' \in A' \cup B'}, (R_e)_{e \in E''})$.

It is obvious by the construction that the instance is (d_A, d_B)-bounded degree. Before we prove the completeness and soundness of the reduction, let us briefly give probabilistic bounds on the sizes of $|E_1|$ and $|E_1 \setminus E''|$ that will be useful in both cases.

Let G_1 denote (A', B', E_1); furthermore, let X_e denote the indicator variable whether the edge e is included in E_1. Let $\mathcal{E}_1$ denote the event that $|E_1| \in [(1-2\lambda)n_E, n_E]$. First, we have $\mathbb{E}[|E_1|] = p|E'| = (1-\lambda)n_E$. Meanwhile, $\mathrm{Var}[|E_1|] = p(1-p)|E'| \leq p|E'| \leq n_E$. As a result, by Chebyshev's inequality, we have

$$\Pr[\neg\mathcal{E}_1] \leq \frac{n_E}{\lambda^2 n_E^2} \leq \frac{1}{\lambda^2 n_E} \leq 0.01,$$

where the last inequality is due to our choice of on d_0.

Let the event $\mathcal{E}_2$ denote the event that $|E_1 \setminus E''| < \lambda|E_1|$. We have

$$\begin{aligned}&\mathbb{E}[|E_1 \setminus E''|] \\ &\leq \sum_{a \in A'} \mathbb{E}\left[\deg_{G_1}(a) - \mathrm{clip}_{d_A}(\deg_{G_1}(a))\right] + \sum_{b \in B'} \mathbb{E}\left[\deg_{G_1}(b) - \mathrm{clip}_{d_B}(\deg_{G_1}(b))\right].\end{aligned}$$

Observe that $\deg_{G_1}(a)$ (resp. $\deg_{G_1}(b)$) is a sum of $d_A C$ (resp. $d_B C$) i.i.d. random variables with mean p. As such, we may apply Lemma 6 to arrive at

$$\begin{aligned}\mathbb{E}[|E_1 \setminus E''|] \leq |A'| \cdot \frac{(d_A C p)^2}{(d_A - d_A C p)^2} + |B'| \cdot \frac{(d_B C p)^2}{(d_B - d_B C p)^2} &\leq (|A'| + |B'|) \cdot \frac{1}{\lambda^2} \\ &\leq \frac{2n_E}{d_0 \lambda^2} \\ &\leq 0.01\lambda|E'|,\end{aligned}$$

where the last inequality follows from our choice of d_0. By Markov's inequality, we thus have

$$\Pr[\neg\mathcal{E}_2] \leq 0.01.$$

Completeness. Henceforth, for any assignment ψ, we use the notation $E'(\psi)$ (resp. $E_1(\psi)$, $E''(\psi)$) to denote the set of edges in E' (resp. E_1, E'') satisfied by ψ.

Let ψ^* be such that $\mathrm{val}_{\Pi'}(\psi^*) = \mathrm{val}(\Pi')$. Let $\mathcal{E}_3$ denote the event $E_1(\psi^*) \geq (\mathrm{val}(\Pi') - 2\lambda) \cdot n_E$. Notice that $E_1(\psi^*)$ is exactly a subset of $E'(\psi^*)$ where each satisfied edge is included with probability p. As a result, we have $\mathbb{E}[|E_1(\psi^*)|] = p \cdot |E'(\psi^*)| = p \cdot |E| \cdot \mathrm{val}(\Pi') \geq n_E \cdot (\mathrm{val}(\Pi') - \lambda)$. Meanwhile, we have $\mathrm{Var}[|E_1(\psi^*)|] = p(1-p)|E'(\psi^*)| \leq n_E \cdot \mathrm{val}(\Pi')$. Thus, by Chebyshev's inequality, we have

$$\Pr[\neg\mathcal{E}_3] \leq \frac{n_E \cdot \mathrm{val}(\Pi')}{(\lambda n_E)^2} \leq \frac{1}{\lambda^2 n_E} \leq 0.01.$$

Thus, by union bound, we have $\Pr[\mathcal{E}_1 \wedge \mathcal{E}_2 \wedge \mathcal{E}_3] \geq 0.97$. When $\mathcal{E}_1, \mathcal{E}_2, \mathcal{E}_3$ all occur, we have

$$\begin{aligned}\mathrm{val}(\Pi'') \geq \mathrm{val}_{\Pi''}(\psi^*) = \frac{|E''(\psi^*)|}{|E''|} \geq \frac{|E_1(\psi^*)| - |E_1 \setminus E''|}{|E_1|} &\geq \frac{(\mathrm{val}(\Pi') - 2\lambda) \cdot n_E - \lambda \cdot n_E}{n_E} \\ &\geq \mathrm{val}(\Pi') - \delta.\end{aligned}$$

Soundness. Assume that $\mathrm{val}(\Pi') \leq \frac{1}{R^\nu}$ and recall that $\chi = \frac{1}{\nu - 2\lambda}\left(\frac{1}{d_A} + \frac{t}{d_B}\right)$ is slightly smaller than the target soundness. Consider any assignment ψ. Let $\mathcal{E}_\psi$ denote the event that $|E_1(\psi)| \leq \chi \cdot n_E$. Again, notice that $E_1(\psi)$ is exactly a subset of $E'(\psi)$ where each satisfied edge is included with probability p. Note also that $|E'(\psi)| \leq \mathrm{val}(\Pi') \cdot |E'|$, implying that $\mathbb{E}[|E_1(\psi)|] = p \cdot |E'(\psi)| \leq \mathrm{val}(\Pi') \cdot n_E$. Thus, we may apply Theorem 5 (with $\theta = \chi \cdot n_E$), to arrive at

$$\begin{aligned}\Pr[\neg\mathcal{E}_\psi] \leq \exp(\chi \cdot n_E - \mathrm{val}(\Pi') \cdot n_E) \cdot \left(\frac{\mathrm{val}(\Pi') \cdot n_E}{\chi \cdot n_E}\right)^{\chi \cdot n_E} &\leq \left(\frac{e}{\chi \cdot R^\nu}\right)^{\chi \cdot n_E} \\ &\leq R^{-(\nu - \lambda) \cdot \chi \cdot n_E}\end{aligned}$$

where the third inequality is from our assumption that $R \geq R_0$.

Therefore, by taking the union bound over all (at most $R^{|A'|} \cdot R^{t|B'|}$) assignments ψ, we have

$$\begin{aligned}\Pr\left[\bigvee_\psi \neg\mathcal{E}_\psi\right] \leq R^{|A'|} \cdot R^{t|B'|} \cdot R^{-(\nu-\lambda)\cdot\chi\cdot n_E} &= R^{\frac{n_E}{d_A} + t \cdot \frac{n_E}{d_B} - (\nu - \lambda)\cdot\chi\cdot n_E} \\ &= (R^{n_E})^{\frac{1}{d_A} + \frac{t}{d_B} - (\nu-\lambda)\cdot\chi} \leq 0.01.\end{aligned}$$

Thus, by the union bound, $\mathcal{E}_1, \mathcal{E}_2$ and $\mathcal{E}_\psi$ for all assignments ψ occur simultanously with probability at least 0.97. When this is the case, we have

$$\mathrm{val}(\Pi'') = \max_\psi \frac{|E''(\psi)|}{|E''|} \leq \frac{|E_1(\psi)|}{|E_1| - |E_1 \setminus E''|} \leq \frac{\chi \cdot n_E}{(1-\lambda)n_E} \leq \frac{1}{\nu - \delta}\left(\frac{1}{d_A} + \frac{t}{d_B}\right),$$

which concludes our proof. ◀

3.3 Putting Things Together: Proof of Theorem 1 and Theorem 2

Our main theorems (Theorems 1 and 2) now follow easily from plugging in the reduction above to the existing large-gap hardness results for 2-CSPs (Theorems 8 and 9) with appropriate parameters.

Proof of Theorem 1. We will reduce from Theorem 9 with $\zeta = 0.01\epsilon$. Let Π be any bipartite (d_1, d_2)-biregular 2-CSP instance with alphabet size R. We first apply Lemma 10 with $c_1 = c_2 = d$ to arrive at a (dd_1d_2, dd_1d_2)-biregular 2-CSP instance Π' with the same alphabet size such that $\text{val}(\Pi') = \text{val}(\Pi)$. We then apply the reduction from Theorem 11 with $d_A = d_B = d$, $t = 1, \delta = 0.01\epsilon, \nu = 1 - \delta$ to arrive at a d-degree-bounded 2-CSP instance Π''. When d is sufficiently large (depending on ϵ only) and R is sufficiently large (depending on d, ϵ), with probability 2/3, we have

- If $\text{val}(\Pi) \geq 1 - \zeta$, then $\text{val}(\Pi'') \geq 1 - \zeta - \delta = 1 - 0.02\epsilon$.
- If $\text{val}(\Pi'') \leq \frac{1}{R^{1-\delta}}$, then $\text{val}(\Pi'') \leq \frac{1}{1-2\delta}\left(\frac{1}{d} + \frac{1}{d}\right) = \frac{1}{1-0.02\epsilon}\left(\frac{2}{d}\right)$.

Note that the ratio between the two cases are larger than $d(1/2 - \epsilon)$. Thus, if there is a polynomial-time $d(1/2 - \epsilon)$-approximation algorithm for 2-CSP on d-bounded-degree graphs, we can distinguish the two cases in randomized polynomial time (with two-sided error). Assuming UGC, from Theorem 9, this implies that NP = BPP. ◀

Proof of Theorem 2. We will reduce from Theorem 8 with $\zeta = 0.01\epsilon$. Let Π be any bipartite (d_1, d_2)-biregular 2-CSP instance with left alphabet size R and right alphabet size $R^{1/2}$. We first apply Lemma 10 with $c_1 = c_2 = d$ to arrive at a (dd_1d_2, dd_1d_2)-biregular 2-CSP instance Π' with the same left and right alphabet sizes such that $\text{val}(\Pi') = \text{val}(\Pi)$. We then apply the reduction from Theorem 11 with $d_A = d_B = d$, $t = 1/2, \delta = 0.01\epsilon, \nu = 1/2 - \delta$ to arrive at a d-degree-bounded 2-CSP instance Π''. When d is sufficiently large (depending on ϵ only) and R is sufficiently large (depending on d, ϵ), with probability 2/3, we have

- If $\text{val}(\Pi) \geq 1 - \zeta$, then $\text{val}(\Pi'') \geq 1 - \zeta - \delta = 1 - 0.02\epsilon$.
- If $\text{val}(\Pi'') \leq \frac{1}{R^{1/2-\delta}}$, then $\text{val}(\Pi'') \leq \frac{1}{1/2-2\delta}\left(\frac{1}{d} + \frac{1/2}{d}\right) = \frac{1}{1-0.04\epsilon}\left(\frac{3}{d}\right)$.

Note that the ratio between the two cases are larger than $d(1/3 - \epsilon)$. Thus, if there is a polynomial-time $d(1/3 - \epsilon)$-approximation algorithm for 2-CSP on d-bounded-degree graphs, we can distinguish the two cases in randomized polynomial time (with two-sided error). From Theorem 8, this implies that NP = BPP. ◀

4 Hardness of Maximum Independent Set in k-Claw-Free Graphs

We next move on to prove hardness of Maximum Independent Set in k-claw-free graphs. To do so, let us first recall the reduction from Max 2-CSP with bounded degree from [15]. As touched on briefly in the introduction, the version we state below is actually more flexible than that in [15] as it allows the degree bounds of the two sides to be different.

▶ **Lemma 12** ([15]). *There is a polynomial-time reduction that takes in a (d_A, d_B)-bounded degree bipartite 2-CSP instance $\Pi = (G = (A, B, E), (\Sigma_v)_{v \in A \cup B}, (R_e)_{e \in E})$ and produces a $(d_A + d_B)$-claw-free graph $G^* = (V^*, E^*)$ such that* $\text{indep}(G^*) = \text{val}(\Pi) \cdot |E|$.

Proof. The reduction is exactly as the so-called "FGLSS graph" [16]:

- For every edge $e = (a, b) \in E$ and every $(\sigma_a, \sigma_b) \in R_e$, create a vertex $(a, b, \sigma_a, \sigma_b)$ in V^*.
- Create an edge between $(a, b, \sigma_a, \sigma_b)$ and $(a', b', \sigma'_{a'}, \sigma'_{b'})$ iff they are inconsistent, i.e. there exists $v' \in \{a, b\} \cap \{a', b'\}$ such that $\sigma_v \neq \sigma'_{v'}$.

A standard argument shows that $\mathrm{indep}(G^*) = \mathrm{val}(\Pi) \cdot |E|$. To see that it is $(d_A + d_B)$-claw-free, consider any vertex $(a, b, \sigma_a, \sigma_b) \in V^*$ and $d_A + d_B$ of its neighbors in G^*. Let $E_0 \subseteq E$ denote the set of edges (in the constraint graph G) adjacent to a or b (or both). By the degree constraint, $|E_0| \leq d_A + d_B - 1$. Thus, by the pigeonhole principle, at least two neighbors of the $d_A + d_B$ neighbors correspond to the same edge in E_0; this means that there must be an edge between these two vertices in G^*. Thus, the graph G^* is $(d_A + d_B)$-claw-free. ◀

Theorem 3 is now an immediate consequence of Theorem 1 and Lemma 12.

Proof of Theorem 3. Note that the instance produced in Theorem 1 is a (d, d)-bounded-degree instance. Thus, by setting $d = \lfloor k/2 \rfloor$ and plugging it into the reduction in Lemma 12, we arrive at the claimed hardness result. ◀

For Theorem 4, we need to work harder to optimize the hardness of approximation factor. Specifically, we set $d_A \approx \sqrt{2} \cdot d_B$, as formalized below.

Proof of Theorem 4. Let $q_1, q_2 \in \mathbb{N}$ be integers such that $\left|\frac{q_1}{q_2} - \sqrt{2}\right|$ and $\left|\frac{q_2}{q_1} - \frac{1}{\sqrt{2}}\right| \leq 0.01\epsilon$. Let $d_A = \lfloor \frac{kq_1}{q_1+q_2} \rfloor$ and $d_B = \lfloor \frac{kq_2}{q_1+q_2} \rfloor$. Note that $d_A + d_B \leq k$.

We will reduce from Theorem 8 with $\zeta = 0.01\epsilon$. Let Π be any bipartite (d_1, d_2)-biregular 2-CSP instance with left alphabet size R and right alphabet size $R^{1/2}$. We first apply Lemma 10 with $c_1 = d_A, c_2 = d_B$ to arrive at a $(d_A d_1 d_2, d_B d_1 d_2)$-biregular 2-CSP instance Π' with the same left and right alphabet sizes such that $\mathrm{val}(\Pi') = \mathrm{val}(\Pi)$. We then apply the reduction from Theorem 11 with d_A, d_B as specified above, $t = 1/2, \delta = 0.01\epsilon, \nu = 1/2 - \delta$ to arrive at a d-degree-bounded 2-CSP instance Π''. Finally, we apply reduction in Lemma 12 on Π'' to arrive at the graph G^*. By Lemma 12, G^* is k-claw-free. Furthermore, when k is sufficiently large (depending on ϵ only) and R is sufficiently large (depending on k, ϵ), with probability $2/3$, we have

- If $\mathrm{val}(\Pi) \geq 1 - \zeta$, then $\mathrm{indep}(G^*) = |E''| \cdot \mathrm{val}(\Pi'') \geq |E''| \cdot (1 - \zeta - \delta) = |E''| \cdot (1 - 0.02\epsilon)$.
- If $\mathrm{val}(\Pi'') \leq \frac{1}{R^{1/2-\delta}}$, then

$$\begin{aligned}
\mathrm{indep}(G^*) = |E''| \cdot \mathrm{val}(\Pi'') &\leq |E''| \cdot \frac{1}{1/2 - 2\delta}\left(\frac{1}{d_A} + \frac{1/2}{d_B}\right) \\
&= \frac{1}{1/2 - 2\delta}\left(\frac{1}{\frac{kq_1}{q_1+q_2} - 1} + \frac{1/2}{\frac{kq_2}{q_1+q_2} - 1}\right) \cdot |E''| \\
&\leq \frac{1}{1/2 - 2\delta} \cdot \frac{1}{1-\delta}\left(\frac{q_1+q_2}{kq_1} + \frac{(q_1+q_2)/2}{kq_2}\right) \cdot |E''| \\
&\leq \frac{1}{1/2 - 2\delta} \cdot \frac{1}{1-\delta} \cdot \frac{1}{k}\left(\frac{3}{2} + \frac{q_2}{q_1} + \frac{q_1}{2q_2}\right) \cdot |E''| \\
&\leq \frac{1}{1/2 - 2\delta} \cdot \frac{1}{1-\delta} \cdot \frac{1}{k}\left(\frac{3}{2} + \sqrt{2} + 2\delta\right) \cdot |E''| \\
&\leq \frac{1}{1 - 4\delta} \cdot \frac{1}{k} \cdot (3 + 2\sqrt{2} + 4\delta) \cdot |E''|,
\end{aligned}$$

  where the second inequality holds when we assume that k is sufficiently large and the second-to-last inequality is from our choice of q_1, q_2.

Note that the ratio between the two cases are larger than $k\left(\frac{1}{3+2\sqrt{2}} - \epsilon\right)$. Thus, if there is a polynomial-time $k\left(\frac{1}{3+2\sqrt{2}} - \epsilon\right)$-approximation algorithm for maximum independent set on k-claw-free graphs, we can distinguish the two cases in randomized polynomial time (with two-sided error). Assuming UGC, from Theorem 9, this implies that NP = BPP. ◀

5 Conclusion and Open Questions

In this paper, we prove hardness of approximation results for Max 2-CSP with bounded degree. Our UG-hardness is nearly tight as the maximum degree goes to ∞. Using this, we also give hardness for Maximum Independent Set on k-claw-free graphs whose inapproximation ratio is within a factor of 2 of optimal for any sufficiently large k. It remains an intriguing open question to close this latter gap. Furthermore, since our reductions are randomized, it would be interesting to derandomized them. Finally, one of our motivations to study Maximum Independent Set on k-claw-free graphs is to understand k-Set Packing. However, we are unable to obtain $\Omega(k)$ factor hardness of approximation of the latter using the reductions in this paper. As stated earlier, the best (NP-)hardness of approximation for k-Set Packing remains $\Omega(k/\log k)$ [23] and it would be interesting to close (or at least decrease) this $O(\log k)$ gap between the upper and lower bounds.

References

1 Per Austrin, Subhash Khot, and Muli Safra. Inapproximability of vertex cover and independent set in bounded degree graphs. *Theory Comput.*, 7(1):27–43, 2011. doi:10.4086/toc.2011.v007a003.

2 Nikhil Bansal, Anupam Gupta, and Guru Guruganesh. On the Lovász theta function for independent sets in sparse graphs. *SIAM J. Comput.*, 47(3):1039–1055, 2018. doi:10.1137/15M1051002.

3 Boaz Barak, Ankur Moitra, Ryan O'Donnell, Prasad Raghavendra, Oded Regev, David Steurer, Luca Trevisan, Aravindan Vijayaraghavan, David Witmer, and John Wright. Beating the random assignment on constraint satisfaction problems of bounded degree. In *APPROX*, pages 110–123, 2015. doi:10.4230/LIPIcs.APPROX-RANDOM.2015.110.

4 Piotr Berman. A $d/2$ approximation for maximum weight independent set in d-claw free graphs. *Nord. J. Comput.*, 7(3):178–184, 2000.

5 Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From gap-exponential time hypothesis to fixed parameter tractable inapproximability: Clique, dominating set, and more. *SIAM J. Comput.*, 49(4):772–810, 2020. doi:10.1137/18M1166869.

6 Parinya Chalermsook, Ameet Gadekar, Kamyar Khodamoradi, and Joachim Spoerhase. Independent set in k-claw-free graphs: Conditional χ-boundedness and the power of LP/SDP relaxations, 2023. arXiv:2308.16033.

7 Siu On Chan. Approximation resistance from pairwise-independent subgroups. *J. ACM*, 63(3):27:1–27:32, 2016. doi:10.1145/2873054.

8 Maria Chudnovsky and Paul Seymour. Claw-free graphs VI. colouring. *Journal of Combinatorial Theory, Series B*, 100(6):560–572, 2010. doi:10.1016/j.jctb.2010.04.005.

9 Marek Cygan. Improved approximation for 3-dimensional matching via bounded pathwidth local search. In *FOCS*, pages 509–518, 2013. doi:10.1109/FOCS.2013.61.

10 Marek Cygan, Fedor V. Fomin, Lukasz Kowalik, Daniel Lokshtanov, Dániel Marx, Marcin Pilipczuk, Michal Pilipczuk, and Saket Saurabh. *Parameterized Algorithms*. Springer, 2015. doi:10.1007/978-3-319-21275-3.

11 Marek Cygan, Fabrizio Grandoni, and Monaldo Mastrolilli. How to sell hyperedges: The hypermatching assignment problem. In *SODA*, pages 342–351, 2013. doi:10.1137/1.9781611973105.25.

12 Michael Dinitz, Guy Kortsarz, and Ran Raz. Label cover instances with large girth and the hardness of approximating basic k-spanner. *ACM Trans. Algorithms*, 12(2):25:1–25:16, 2016. doi:10.1145/2818375.

13 Irit Dinur. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. *Electron. Colloquium Comput. Complex.*, TR16-128, 2016. arXiv:TR16-128.

14 Irit Dinur and Pasin Manurangsi. ETH-hardness of approximating 2-CSPs and directed steiner network. In *ITCS*, pages 36:1–36:20, 2018. doi:10.4230/LIPIcs.ITCS.2018.36.

15 Pavel Dvorák, Andreas Emil Feldmann, Ashutosh Rai, and Pawel Rzazewski. Parameterized inapproximability of independent set in H-free graphs. *Algorithmica*, 85(4):902–928, 2023. doi:10.1007/s00453-022-01052-5.

16 Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. Interactive proofs and the hardness of approximating cliques. *J. ACM*, 43(2):268–292, 1996. doi:10.1145/226643.226652.

17 Joel Friedman. *A proof of Alon's second eigenvalue conjecture and related problems*. American Mathematical Soc., 2008.

18 Oded Goldreich and Madhu Sudan. Locally testable codes and pcps of almost-linear length. *J. ACM*, 53(4):558–655, 2006. doi:10.1145/1162349.1162351.

19 Martin Grötschel, László Lovász, and Alexander Schrijver. The ellipsoid method and its consequences in combinatorial optimization. *Combinatorica*, 1:169–197, 1981.

20 Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *FOCS*, pages 627–636, 1996. doi:10.1109/SFCS.1996.548522.

21 Johan Håstad. On bounded occurrence constraint satisfaction. *Inf. Process. Lett.*, 74(1-2):1–6, 2000. doi:10.1016/S0020-0190(00)00032-6.

22 Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001. doi:10.1145/502090.502098.

23 Elad Hazan, Shmuel Safra, and Oded Schwartz. On the complexity of approximating k-set packing. *Comput. Complex.*, 15(1):20–39, 2006. doi:10.1007/s00037-006-0205-6.

24 Subhash Khot. On the power of unique 2-prover 1-round games. In *CCC*, page 25, 2002. doi:10.1109/CCC.2002.1004334.

25 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.

26 Guy Kindler, Alexandra Kolla, and Luca Trevisan. Approximation of non-boolean 2CSP. In *SODA*, 2016. doi:10.1137/1.9781611974331.ch117.

27 Bundit Laekhanukit. Parameters of two-prover-one-round game and the hardness of connectivity problems. In *SODA*, pages 1626–1643, 2014. doi:10.1137/1.9781611973402.118.

28 Pasin Manurangsi. A note on degree vs gap of min-rep label cover and improved inapproximability for connectivity problems. *Inf. Process. Lett.*, 145:24–29, 2019. doi:10.1016/j.ipl.2018.08.007.

29 Pasin Manurangsi and Prasad Raghavendra. A birthday repetition theorem and complexity of approximating dense CSPs. In *ICALP*, pages 78:1–78:15, 2017. doi:10.4230/LIPIcs.ICALP.2017.78.

30 Pasin Manurangsi, Aviad Rubinstein, and Tselil Schramm. The Strongish Planted Clique Hypothesis and its consequences. In *ITCS*, pages 10:1–10:21, 2021. doi:10.4230/LIPIcs.ITCS.2021.10.

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

32 Meike Neuwohner. An improved approximation algorithm for the maximum weight independent set problem in d-claw free graphs. In *STACS*, pages 53:1–53:20, 2021. doi:10.4230/LIPIcs.STACS.2021.53.

33 Meike Neuwohner. Passing the limits of pure local search for weighted k-set packing. In *SODA*, pages 1090–1137, 2023. doi:10.1137/1.9781611977554.ch41.

34 Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *STOC*, pages 245–254, 2008. doi:10.1145/1374376.1374414.

35 Alexander Schrijver. *Combinatorial optimization: polyhedra and efficiency*, volume 24. Springer, 2003.

36 Maxim Sviridenko and Justin Ward. Large neighborhood local search for the maximum set packing problem. In *ICALP*, pages 792–803, 2013. doi:10.1007/978-3-642-39206-1_67.
37 Theophile Thiery and Justin Ward. An improved approximation for maximum weighted k-set packing. In *SODA*, pages 1138–1162, 2023. doi:10.1137/1.9781611977554.ch42.
38 Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *STOC*, pages 453–461, 2001. doi:10.1145/380752.380839.
39 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory Comput.*, 3(1):103–128, 2007. doi:10.4086/toc.2007.v003a006.

A UGC-Hardness of 2-CSP with Almost Perfect Completeness

In this section, we prove Theorem 9. It follows from a standard technique proving hardness of CSPs assuming the Unique Games Conjecture [25].

Proof of Theorem 9. For given $R \in \mathbb{N}$, we will construct a *predicate* $P \subseteq [R] \times [R]$, and consider CSP($P$) whose instance $\Pi = (G = (V, E), (\Sigma_v)_{v \in V}, (R_e)_{e \in E})$ must satisfy $\Sigma_v = [R]$ for every $v \in V$ and $R_e = \{(x, y) \in [R]^2 : (x \oplus t_{e,u}, y \oplus t_{e,v}) \in P\}$ for some $t_{e,u}, t_{e,v} \in [R]$ for every $e = (u, v) \in E$. (For $x, y \in [R]$, we define $x \oplus y$ to be $x + y$ if it is at most R and $x + y - R$ otherwise.)

The standard technique of proving the hardness of CSP(P) due to Khot et al. [25] shows that it suffices to consider the *dictatorship test*. It is determined by a distribution μ supported on $[R] \times [R]$. For every $L \in \mathbb{N}$, it yields the following test that decides whether a given function $F : [R]^L \to [R]$ is a *dictator* or not.

- For each $i \in [L]$, sample $(x_i, y_i) \in [R]^2$ from μ, independently from other i's.
- Accept if $(F(x), F(y)) \in P$.

Note that if F is a dictator function (i.e., $F(x) = x_i$ for some $i \in [L]$), then the above test accepts with probability exactly $\Pr_{(x,y)\sim\mu}[(x, y) \in P] =: c$, known as the completeness of the test. Let s be (an upper bound of) the soundness of this test; there exist $\tau > 0$ and $d \in \mathbb{N}$ such that any function $F : [R]^L \to [R]$, which (1) is *balanced* (i.e., $|F^{-1}(i)| = R^{L-1}$ for all $i \in [R]$) and (2) has the *maximum degree-d influence* (defined in Appendix A.2) at most τ, passes the test with probability at most s. Khot et al. [25] shows that a dictatorship test with some c and s immediately yields the hardness of CSP(P) with almost the same completeness and soundness.

▶ **Theorem 13** ([25]). *Given $P \subseteq [R]^2$, let μ be a distribution over $[R]^2$ that yields a dictatorship test with completeness c and soundness s. Then, for any $\zeta > 0$, assuming the Unique Games Conjecture, it is NP-hard, given a regular CSP(P) instance Π, to distinguish between the following two cases:*
- *(Yes Case)* $\mathrm{val}(\Pi) \geq c - \zeta$.
- *(No Case)* $\mathrm{val}(\Pi) \leq s + \zeta$.

Though the above theorem does not guarantee that a given instance Π with the underlying graph $G = (V, E)$ is bipartite, one can easily convert it to a bipartite instance Π' by creating two vertex sets V_1 and V_2 which are disjoint copies of V and replace each constraint $e = (u, v)$ with $R_e \subseteq [R]^2$ with two constraints (u_1, v_2) and (u_2, v_1) with the same R_e, where u_i, v_i denote the copy of u, v in V_i. The completeness of the new instance is at least the completeness of the original instance, and the soundness of the new instance is at most twice of the soundness of the original instance.

Therefore, the rest of the section is devoted to constructing P and μ such that μ is supported by P (so that $c = 1$) while $s = O(\log^2 R/R)$.

A.1 Predicate and Completeness

Given $R \in \mathbb{N}$, let $t \in \mathbb{N}$ be a parameter to be determined later, and let $H = (V_H, E_H)$ be a t-regular graph with $V_H = [R]$ such that the second largest eigenvalue of the normalized adjacency matrix is $O(1/\sqrt{t})$ [17]. The predicate $P \subseteq [R]^2$ is defined such that $(i, j) \in P$ if and only if $(i, j) \in E_H$. This ensures that $|P| = tR = 2|E_H|$. Then, our distribution μ is simply the uniform distribution over P. By definition, the completeness value is $\Pr_{(x,y)\sim\mu}[(x, y) \in P] = 1$.

A.2 Soundness via Fourier analysis

To (formally define and) analyze the soundness of the test, we use the following standard tools from Gaussian bounds for correlated functions from Mossel [31]. We define the correlation between two correlated spaces below.

▶ **Definition 14.** *Given a distribution μ on $\Omega_1 \times \Omega_2$, the correlation $\rho(\Omega_1, \Omega_2; \mu)$ is defined as*

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

In our case, $\rho := \rho([R], [R]; \mu)$ is exactly the second largest eigenvalue of the normalized adjacency matrix of H, which is $O(1/\sqrt{t})$.

▶ **Definition 15** ([31]). *For any function $f : [R]^L \to \mathbb{R}$, the* Efron-Stein *decomposition is given by*

$$f(y) = \sum_{S\subseteq[L]} f_S(y)$$

where the functions f_S satisfy
- *f_S only depends on y_S, the restriction of y to the coordinates of S.*
- *For all $S \not\subseteq S'$ and all $z_{S'}$, $\mathbb{E}_y[f_S(y)|y_{S'} = z_{S'}] = 0$.*

Based on the Efron-Stein decomposition, we can define (low-degree) influences of a function. For a function $f : [R]^L \to \mathbb{R}$ and $p \geq 1$, let $\|f\|_p := \mathbb{E}[|f(y)|^p]^{1/p}$.

▶ **Definition 16** ([31]). *For any function $f : [R]^L \to \mathbb{R}$, its ith* influence *is defined as*

$$\mathsf{Inf}_i(f) := \sum_{S:i\in S} \|f_S\|_2^2.$$

Its ith degree-d influence *is defined as*

$$\mathsf{Inf}_i^{\leq d}(f) := \sum_{S:i\in S,|S|\leq d} \|f_S\|_2^2,$$

Given a discrete-valued function $F : [R]^L \to [R]$, for every $i \in [R]$, we let $F_i : [R]^L \to \{0, 1\}$ such that $F_i(x) = 1$ if $F(x) = i$ and 0 otherwise. We say that F has the maximum degree-d influence at most τ if $\mathsf{Inf}_j^{\leq d}(F_i) \leq \tau$ for every $i \in [R]$ and $j \in [L]$.

For $a, b \in [0, 1]$ and $\sigma \in [0, 1]$, let $\Gamma_\sigma(a, b) := \Pr[g_1 \leq \Phi^{-1}(a), g_2 \leq \Phi^{-1}(b)]$ where g_1, g_2 are σ-correlated standard Gaussian variables and Φ denotes the cumulative density function of a standard Gaussian. (E.g., $\Gamma_\sigma(a, 1) = a$ for any σ and $\Gamma_0(a, b) = ab$.) We crucially use the following invariance principle applied to our dictatorship test.

▶ **Theorem 17** ([31])**.** *For any $\epsilon > 0$ there exist $d \in \mathbb{N}$ and $\tau > 0$ such that the following is true. Let $f, g : [R]^L \to [0,1]$. If $\min(\mathsf{Inf}_i^{\le d}[f], \mathsf{Inf}_i^{\le d}[g]) \le \tau$ for every $i \in [L]$,*

$$\mathbb{E}_{(x,y)\sim\mu^{\otimes L}}[f(x)g(y)] \le \Gamma_\rho(\mathbb{E}[f(x)], \mathbb{E}[g(y)]) + \epsilon.$$

We will use the following upper bound on $\Gamma_\rho(\alpha, \alpha)$.

▶ **Lemma 18** (Corollary 3 of [25])**.** *For any $R \ge 1$ and $\rho \in (0, 1/20)$,*

$$\Gamma_\rho(1/R, 1/R) \le (1/R)^{1+\frac{1-\rho}{1+\rho}} \le (1/R)^{2-2\rho}.$$

Fix $\epsilon = 1/R^3$ to get d and τ from Theorem 17. Then, for any $F : [R]^L \to [R]$ that is balanced (i.e., $|F^{-1}(i)| = R^{L-1}$ for every $i \in [R]$) and has the maximum degree-d influence at most τ, the probability that the dictatorship test accepts is

$$\sum_{(i,j)\in P} (\mathbb{E}_{(x,y)\sim\mu^{\otimes L}}[F_i(x)F_j(y)] + \epsilon) \le tR \cdot (\Gamma_\rho(1/R, 1/R) + \epsilon) \le t \cdot (1/R)^{1-2\rho} + 1/R.$$

Recalling $\rho = O(1/\sqrt{t})$ and setting $t = \Theta(\log^2 R)$ ensure that $t(1/R)^{1-2\rho} \le O(\log^2 R/R)$, so the soundness is at most $O(\log^2 R/R)$. ◀

B Approximation Algorithm

In this section, we give a $(\frac{d+1}{2})$-approximation full algorithm for any d-bounded-degree 2-CSP. Before we proceed to the algorithm, let us note that in the case where the instance is fully satisfiable, there is a simple algorithm: Just take any spanning forest of the constraint graph and then use a dynamic programming algorithm to find an assignment that satisfies all the edges in the spanning forest! This algorithm does not work in the general case since it is possible that this spanning forest has a small value. To overcome this, below we sample the spanning forest from an appropriate distribution, allowing us to maintain the same approximation ratio.

▶ **Theorem 19.** *There is a polynomial-time $(\frac{d+1}{2})$-approximation algorithm for every d-bounded-degree 2-CSP.*

Proof. Let $\Pi = (G = (V, E), (\Sigma_v)_{v\in V}, (R_e)_{e\in E})$ be an instance of 2-CSP where the maximum degree of G is at most d.

Let $x \in \mathbb{R}^E$ be such that $x_e = 2/(d+1)$ for every $e \in E$. We claim that x is inside the graphic matroid polytope induced by G; for any $S \subseteq V$, if $|S| \in [2, d+1]$, $x(E(S)) \le \frac{|S|(|S|-1)}{2}(\frac{2}{d+1}) \le |S| - 1$, and if $|S| > d+1$, $x(E(S)) \le \frac{d|S|}{2} \cdot (\frac{2}{d+1}) \le |S| - 1$.

Therefore, x can be written as a convex combination of (the indicator vectors of) forests in G [35], which implies that there exists a distribution $\mathcal{T}$ of forests such that for a random forest $T \sim \mathcal{T}$, for every $e \in E$, $\Pr[e \in T] \ge \frac{2}{d+1}$. Then, using dynamic programming, one can optimally solve the subinstance of Π induced by T; for each connected component T' of T (which is a tree), root it at an arbitrary vertex, and for each node $v \in T'$ and $\sigma \in \Sigma_v$, let $A(v, \sigma)$ be the the optimal value of the CSP induced by the subtree of T' rooted at v when the variable v is assigned label σ. One can compute $A(v, \sigma)$ in a bottom-up fashion using dynamic programming.

Since $\Pr[e \in T] \ge \frac{2}{d+1}$ for every $e \in E$, the expected optimal number of satisfied constraints of the CSP instance induced by T is at least $\frac{2}{d+1} \cdot |E| \cdot \mathrm{val}(\Pi)$. Therefore, returning an optimal assignment for a random T yields a $\left(\frac{d+1}{2}\right)$-approximation in expectation. It can be easily derandomized since the integrality of the graphic matroid polytope implies that one can efficiently compute a decomposition of x into a convex combination of $|E| + 1$ forests [19]. ◀

C Parameterized Hardness of Approximation

In this section, we briefly discussed the parameterized hardness of approximation for the Maximum Independent Set in k-claw-free graphs. Recall that an algorithm is said to be fixed parameter tractable (FPT) w.r.t. to parameter q if it runs in time $f(q) \cdot n^{O(1)}$ where f can be any function and n is the input size. We refer interested readers to [10] for more background on the topic.

Similar to [15], we let the parameter be $q = k + \text{indep}(G)$. For this parameter, [15] showed (by reducing from parameterized hardness of Max 2-CSP in [30]) that, assuming the (less standard) Strongish Planted Clique Hypothesis, no FPT algorithm achieves $o(k)$-approximation. Note that this is incomparable to hardness presented in the main body of our paper (Theorems 3 and 4), as such a parameterized hardness result does not rule out e.g. $n^{O(k)}$-time algorithm. (Our main results rule out such algorithms since k there are simply absolute constants.)

Meanwhile, under the (arguably more standard) Gap-ETH assumption[8], [15] only show (via a reduction from parameterized hardness of Max 2-CSP in [14]) that no FPT algorithm achieves $o\left(\frac{k}{2^{(\log k)^{1/2+o(1)}}}\right)$-approximation. Our result here is an improvement of this factor to $o(k)$:

▶ **Theorem 20.** *Assuming Gap-ETH, there is no $f(k) \cdot n^{O(1)}$-time $o(k)$-approximation algorithm for Maximum Independent Set on k-claw-free graphs even when the maximum independent set has size at most k.*

To prove this theorem, we need the following additional notations for 2-CSPs:

- For a 2-CSP instance $\Pi = (G = (V, E), (\Sigma_v)_{v \in V}, (R_e)_{e \in E})$, a *partial assignment* is a tuple $(\psi_v)_{v \in V}$ such that $\psi_v \in \Sigma_v \cup \{\perp\}$. Its *size* is defined as $|\{v : \psi_v \neq \perp\}|$.
- We say that a partial assignment ψ is consistent if, for all $e = (u, v) \in V$ such that $\psi_u, \psi_v \neq \perp$, we have $(\psi_u, \psi_v) \in R_e$.
- Finally, we define $\text{cval}(\Pi)$ to be the maximum size of any consistent partial assignment.

We will use the following hardness result[9]:

▶ **Theorem 21** ([5]). *Assuming Gap-ETH, there is no $f(k) \cdot n^{O(1)}$-time $o(k)$-approximation algorithm for* $\text{cval}(\Pi)$ *with k variables.*

Our main ingredient is the following reduction, which is different than that of [15] and allows us to use cval instead of val for (in)approximation purposes.

▶ **Lemma 22.** *There is a polynomial-time reduction that takes in a 2-CSP instance $\Pi = (G = (V, E), (\Sigma_v)_{v \in V}, \{R_e\}_{e \in E})$ and produces a graph $G' = (V', E')$ such that* $\text{indep}(G') = \text{cval}(\Pi)$. *Moreover, if G has degree at most d, then G' is $(d + 2)$-claw-free.*

Proof. Let G' be the label-extended graph of G. Namely, $V' = \{(v, \sigma_v) \mid v \in V, \sigma_v \in \Sigma_v\}$ and there is an edge between (u, σ_u) and (v, σ_v) in E' iff $(u, v) \in E$ and $(\sigma_u, \sigma_v) \notin R_e$. The claim $\text{cval}(\Pi) = \text{indep}(V')$ is obvious. To see that the graph G' is $(d + 2)$-claw-free, observe that any vertex (u, σ_u) is only neighbors to (v, σ_v) where $v \in N_G[u]$ (the closed-neighbor of G). However, for each fixed v, $\{(v, \sigma_v) \mid \sigma_v \in \Sigma_v\}$ forms a clique. Thus, the largest size of claw that is a subgraph of G' is at most $|N_G[u]| \leq d + 1$. ◀

Plugging in the above lemma to Theorem 21, we immediately arrive at Theorem 20.

[8] Gap Exponential Time Hypothesis (Gap-ETH) [13, 29] states that no $2^{o(n)}$-time algorithm can distinguish between a fully satisfiable 3-SAT instance and one which is not even $(1 - \epsilon)$-satisfiable for some constant $\epsilon > 0$.

[9] In [5], this is stated as the hardness of Clique, but this is exactly the same as 2-CSP with k variables.