Bilu-Linial Stability, Certified Algorithms and the Independent Set Problem

Haris Angelidakis
ETH Zürich, Switzerland

Pranjal Awasthi
Rutgers University, New Brunswick, NJ, USA

Avrim Blum
Toyota Technological Institute at Chicago, IL, USA

Vaggos Chatziafratis
Stanford University, CA, USA

Chen Dan
Carnegie Mellon University, Pittsburgh, PA, USA

Abstract

We study the classic Maximum Independent Set problem under the notion of *stability* introduced by Bilu and Linial (2010): a weighted instance of Independent Set is γ-stable if it has a unique optimal solution that remains the unique optimal solution under multiplicative perturbations of the weights by a factor of at most $\gamma \geq 1$. The goal then is to efficiently recover this "pronounced" optimal solution exactly. In this work, we solve stable instances of Independent Set on several classes of graphs: we improve upon previous results by solving $\widetilde{O}(\Delta/\sqrt{\log \Delta})$-stable instances on graphs of maximum degree Δ, $(k-1)$-stable instances on k-colorable graphs and $(1+\varepsilon)$-stable instances on planar graphs (for any fixed $\varepsilon > 0$), using both combinatorial techniques as well as LPs and the Sherali-Adams hierarchy.

For general graphs, we present a strong lower bound showing that there are no efficient algorithms for $O(n^{\frac{1}{2}-\varepsilon})$-stable instances of Independent Set, assuming the planted clique conjecture. To complement our negative result, we give an algorithm for (εn)-stable instances, for any fixed $\varepsilon > 0$. As a by-product of our techniques, we give algorithms as well as lower bounds for stable instances of Node Multiway Cut (a generalization of Edge Multiway Cut), by exploiting its connections to Vertex Cover. Furthermore, we prove a general structural result showing that the integrality gap of convex relaxations of several maximization problems reduces dramatically on stable instances.

Moreover, we initiate the study of *certified* algorithms for Independent Set. The notion of a γ-certified algorithm was introduced very recently by Makarychev and Makarychev (2018) and it is a class of γ-approximation algorithms that satisfy one crucial property: the solution returned is optimal for a perturbation of the original instance, where perturbations are again multiplicative up to a factor of $\gamma \geq 1$ (hence, such algorithms not only solve γ-stable instances optimally, but also have guarantees even on unstable instances). Here, we obtain Δ-certified algorithms for Independent Set on graphs of maximum degree Δ, and $(1+\varepsilon)$-certified algorithms on planar graphs. Finally, we analyze the algorithm of Berman and Fürer (1994) and prove that it is a $\left(\frac{\Delta+1}{3}+\varepsilon\right)$-certified algorithm for Independent Set on graphs of maximum degree Δ where all weights are equal to 1.

2012 ACM Subject Classification Theory of computation → Packing and covering problems; Theory of computation → Algorithm design techniques

Keywords and phrases Bilu-Linial stability, perturbation resilience, beyond worst-case analysis, Independent Set, Vertex Cover, Multiway Cut

Digital Object Identifier 10.4230/LIPIcs.ESA.2019.25

Related Version A full version of the paper is available at `https://arxiv.org/abs/1810.08414`.

Funding *Haris Angelidakis*: Part of this work was done while the author was a student at TTI-Chicago, and was supported by the National Science Foundation under the grant CCF-1718820.

27th Annual European Symposium on Algorithms (ESA 2019).
Editors: Michael A. Bender, Ola Svensson, and Grzegorz Herman; Article No. 25; pp. 25:1–25:16
Leibniz International Proceedings in Informatics
LIPICS Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Avrim Blum: This work was supported in part by the National Science Foundation under grants CCF-1733556, CCF-1800317, and CCF-1815011.

Acknowledgements We would like to thank Konstantin Makarychev and Yury Makarychev for sharing their manuscript [41], and Yury Makarychev and Mrinalkanti Ghosh for useful discussions.

1 Introduction

The Maximum Independent Set problem (simply MIS from now on) is a central problem in theoretical computer science and has been the subject of extensive research over the last few decades. As a result we now have a thorough understanding of the *worst-case* behavior of the problem. In general graphs, the problem is $n^{1-\varepsilon}$-hard to approximate, assuming that $\mathtt{P} \neq \mathtt{NP}$ [34, 48], and $n/2^{(\log n)^{3/4+\varepsilon}}$-hard to approximate, assuming that $\mathtt{NP} \not\subseteq \mathtt{BPTIME}(2^{(\log n)^{O(1)}})$ [37]. On the positive side, the current best algorithm is due to Feige [27] achieving a $\widetilde{O}(n/\log^3 n)$-approximation (the notation $\widetilde{O}$ hides some $\text{poly}(\log\log n)$ factors). In order to circumvent the strong lower bounds, many works have focused on special classes of graphs, such as bounded-degree graphs (see, e.g., [1, 5, 10, 11, 21, 31, 32, 33]), planar graphs ([7]) etc. In this work, we build upon this long line of research and study MIS under the beyond worst-case framework introduced by Bilu and Linial [17].

In an attempt to capture real-life instances of combinatorial optimization problems, Bilu and Linial proposed a notion of stability, which we now instantiate in the context of MIS (from now on, we will always assume weighted instances of MIS).

▶ **Definition 1** (γ-perturbation [17]). *Let $G = (V, E, w)$, $w : V \to \mathbb{R}_{>0}$, be an instance of MIS. An instance $G' = (V, E, w')$ is a γ-perturbation of G, for some parameter $\gamma \geq 1$, if for every $u \in V$ we have $w_u \leq w'_u \leq \gamma \cdot w_u$.*

▶ **Definition 2** (γ-stability [17]). *Let $G = (V, E, w)$, $w : V \to \mathbb{R}_{>0}$, be an instance of MIS. The instance G is γ-stable, for some parameter $\gamma \geq 1$, if:*
1. *it has a unique maximum independent set I^*,*
2. *every γ-perturbation G' of G has a unique maximum independent set equal to I^*.*

Equivalently, G is γ-stable if it has an independent set I^* such that $w(I^* \setminus S) > \gamma \cdot w(S \setminus I^*)$ for every feasible independent set $S \neq I^*$ (we use the notation $w(Q) := \sum_{u \in Q} w_u$ for $Q \subseteq V$).

This definition of stability is motivated by the empirical observation that in many real-life instances, the optimal solution stands out from the rest of the solution space, and thus is not sensitive to small perturbations of the parameters. This suggests that the optimal solution does not change (structurally) if the parameters of the instance are perturbed (even adversarially). Observe that the smaller the so-called *stability threshold* γ is, the less severe the restrictions imposed on the instance are; for example, $\gamma = 1$ is the case where we only require the optimal solution to be unique. Thus, the main goal in this framework is to recover the optimal solution in polynomial time, for as small $\gamma \geq 1$ as possible. An "optimal" result would translate to γ being $1 + \varepsilon$, for small $\varepsilon > 0$, since assuming uniqueness of the optimal solution is not believed to make the problems easier (see, e.g., [47]), and thus ε is unlikely to be zero. We note that perturbations are scale-invariant, and so it suffices to consider perturbations that only scale up. Moreover, we observe that an algorithm for γ-stable instances of MIS solves γ-stable instances of Minimum Vertex Cover, and vice versa.

Stability was first introduced for Max Cut [17], but the authors note that it naturally extends to other problems, such as MIS, and, moreover, they prove that the greedy algorithm for MIS solves Δ-stable instances on graphs of maximum degree Δ. The work of Bilu and

Linial has inspired numerous works on stable instances of various optimization problems; we give an overview of the literature in the next page.

Prior works on stability have also studied *robust* algorithms [42, 4]; these are algorithms that either output an optimal solution or provide a polynomial-time verifiable certificate that the instance is not γ-stable (see Section 2 for a definition). Motivated by the notion of stability, Makarychev and Makarychev [41] recently introduced an intriguing class of algorithms, namely γ-*certified* algorithms.

▶ **Definition 3** (γ-certified algorithm [41])**.** *An algorithm for* `MIS` *is called γ-certified, for some parameter $\gamma \geq 1$, if for every instance $G = (V, E, w)$, $w : V \to \mathbb{R}_{>0}$, it computes*

1. *a feasible independent set $S \subseteq V$ of G,*
2. *a γ-perturbation $G' = (V, E, w')$ of G such that S is a maximum independent set of G'.*

Equivalently, Condition (2) can be replaced by the following: $\gamma \cdot w(S \setminus I) \geq w(I \setminus S)$ for every independent set I of G.

We highlight that a certified algorithm works for *every* instance; if the instance is γ-stable, then the solution returned is the optimal one, while if it is not stable, the solution is within a γ-factor of optimal. Hence a γ-certified algorithm is also a γ-approximation algorithm.

Motivation. Stability is especially natural for problems where the given objective function may be a proxy for a true goal of identifying a hidden correct solution. For `MIS`, a natural such scenario is applying a machine learning algorithm in the presence of pairwise constraints. Consider, for instance, an algorithm that scans news articles on the web and aims to extract events such as "athlete X won the Olympic gold medal in Y". For each such statement, the algorithm gives a confidence score (e.g., it might be more confident if it saw this listed in a table rather than inferring it from a free-text sentence that the algorithm might have misunderstood). But in addition, the algorithm might also know logical constraints such as "at most one person can win a gold medal in any given event". These logical constraints would then become edges in a graph, and the goal of finding the most likely combination of events would become a `MIS` problem. Stability would be natural to assume in such a setting since the exact confidence weights are somewhat heuristic, and the goal is to recover an underlying ground truth. It is also easy to see the usefulness of a certified algorithm in this setting. Given a certified algorithm that outputs a γ-perturbation, the user of the machine learning algorithm can further test and debug the system by trying to gather evidence for events on which the perturbation puts higher weight.

Related Work. There have been many works on the worst-case complexity of `MIS` and the current best known algorithms give $\widetilde{O}(n/\log^3 n)$-approximation [27], and $\widetilde{O}(\Delta/\log\Delta)$-approximation [31, 33, 36]), where Δ is the maximum degree. The problem has also been studied from the lens of beyond worst-case analysis. For random graphs with a planted independent set, `MIS` is equivalent to the classic planted clique problem. Inspired by semi-random models of [18], Feige and Killian [28] designed SDP-based algorithms for computing large independent sets in semi-random graphs. Finally, there has been work on `MIS` under noise [40, 12].

The notion of Bilu-Linial stability goes beyond random/semi-random models and proposes deterministic conditions that give rise to non worst-case, real-life instances. The study of this notion has led to insights into the complexity of many problems in optimization and machine learning. For `MIS`, Bilu [15] analyzed the greedy algorithm and showed that it recovers the optimal solution for Δ-stable instances of graphs of maximum degree Δ. The same result is also a corollary of a general theorem about the greedy algorithm and p-extendible independence systems proved by Chatziafratis et al. [22]. On the negative side, Angelidakis et al. [4] showed that there is no robust algorithm for $n^{1-\varepsilon}$-stable instances of MIS on general graphs (unbounded degree), assuming that $\mathtt{P} \neq \mathtt{NP}$.

The work of Bilu and Linial has inspired a sequence of works about stable instances of various combinatorial optimization problems. There are now algorithms that solve $O(\sqrt{\log n}\log\log n)$-stable instances of Max Cut [16, 42], $(2-2/k)$-stable instances of Edge Multiway Cut, where k is the number of terminals [42, 4], and 1.8-stable instances of symmetric TSP [45]. There has also been extensive work on stable instances of clustering problems (usually called perturbation-resilient instances) with many positive results for problems such as k-median, k-means, and k-center [6, 9, 8, 4, 25, 23, 26, 30], and more recently on MAP inference [38, 39].

Our results. We explore the notion of stability in the context of MIS and significantly improve our understanding of its behavior on stable instances; we design algorithms for stable instances on different graph classes, and also initiate the study of certified algorithms for MIS. More specifically, we obtain the following results.

- **Planar graphs:** We show that on planar graphs, any constant stability suffices to solve the problem exactly in polynomial time. More precisely, we provide robust and certified algorithms for $(1+\varepsilon)$-stable instances of planar MIS, for any fixed $\varepsilon > 0$. To obtain these results, we utilize the Sherali-Adams hierarchy, demonstrating that hierarchies may be helpful for solving stable instances.
- **Graphs with small chromatic number or bounded degree:** We provide robust algorithms for solving $(k-1)$-stable instances of MIS on k-colorable graphs (where the algorithm does not have access to a k-coloring of the graph) and $(\Delta-1)$-stable instances of MIS on graphs of maximum degree Δ. Both results are based on LPs. For bounded-degree graphs, we then turn to combinatorial techniques and design a (non-robust) algorithm for $\widetilde{O}(\Delta/\sqrt{\log\Delta})$-stable instances; this is the first algorithm that solves $o(\Delta)$-stable instances. Moreover, we show that the standard greedy algorithm is a Δ-certified algorithm for MIS, whereas for unweighted instances, the algorithm of Berman and Fürer (1994) is a $\left(\frac{\Delta+1}{3}+\varepsilon\right)$-certified algorithm.
- **General graphs:** For general graphs, we show that solving $o(\sqrt{n})$-stable instances is hard assuming the hardness of finding maximum cliques in a random graph. To the best of our knowledge, this is only the second case of a lower bound for stable instances of a graph optimization problem that applies to any polynomial-time algorithm and not only to robust algorithms [42, 4] (the first being the lower bound for Max k-Cut [42]). We complement this lower bound by giving an algorithm for (εn)-stable instances of MIS on graphs with n vertices, for any fixed $\varepsilon > 0$.
- **Convex relaxations and stability:** We present a structural result for the integrality gap of convex relaxations of maximization problems on stable instances: if the integrality gap of a relaxation is α, then it is at most $\min\left\{\alpha, 1+\frac{1}{\beta-1}\right\}$ for $(\alpha\beta)$-stable instances, for any $\beta > 1$. This result demonstrates a smooth trade-off between stability and the performance of a convex relaxation, and also implies $(1+\varepsilon)$-estimation algorithms[1] for $O(\alpha/\varepsilon)$-stable instances.

[1] An α-estimation algorithm returns a value that is within a factor of α from the optimum, but not necessarily a corresponding solution that realizes this value.

- **Node Multiway Cut:** We give the first results on stable instances of Node Multiway Cut, a strict generalization of the well-studied (under stability) Edge Multiway Cut problem [42, 4]. In particular, we give a robust algorithm for $(k-1)$-stable instances, where k is the number of terminals, and show that all negative results on stable instances of MIS directly apply to Node Multiway Cut.

Organization of material. Section 2 provides definitions and related facts. Section 3 contains the algorithms for stable instances of MIS on bounded-degree, small chromatic number and planar graphs. Section 4 contains our results for stable instances on general graphs. Section 5 demonstrates how the performance of convex relaxations improves as stability increases. Section 6 contains various certified algorithms for MIS. Due to space constraints, the results for the Node Multiway Cut problem are omitted and can be found in the full version of the paper [3]. The same applies to all proofs that have been omitted.

2 Preliminaries and definitions

Given a γ-stable instance, our goal is to design polynomial-time algorithms that recover the unique optimal solution, for as small $\gamma \geq 1$ as possible. A special class of such algorithms that is of particular interest is the class of robust algorithms, introduced by Makarychev et al. [42].

▶ **Definition 4** (robust algorithm [42]). *Let $G = (V, E, w)$, $w : V \to \mathbb{R}_{>0}$, be an instance of MIS. An algorithm $\mathcal{A}$ is a robust algorithm for γ-stable instances if:*
1. *it always returns the unique optimal solution of G, when G is γ-stable,*
2. *it either returns an optimal solution of G or reports that G is not stable, when G is not γ-stable.*

Note that a robust algorithm is not allowed to err, while a non-robust algorithm is allowed to return a suboptimal solution, if the instance is not γ-stable. We now present a useful lemma about stable instances of MIS that is used in several of our results. From now on, we denote the neighborhood of a vertex u of a graph $G = (V, E)$ as $N(u) = \{v : (u, v) \in E\}$, and the neighborhood of a set $S \subseteq V$ as $N(S) = \{v \in V \setminus S : \exists u \in S \text{ s.t. } (u, v) \in E\}$.

▶ **Lemma 5.** *Let $G = (V, E, w)$ be a γ-stable instance of MIS whose optimal independent set is I^*. Then, for any $v \in I^*$, the induced instance $\widetilde{G} = G[V \setminus (\{v\} \cup N(v))]$ is γ-stable, and its unique maximum independent set is $I^* \setminus \{v\}$.*

Regarding certified algorithms (see Definition 3), it is easy to observe the following.

▶ **Observation 6** ([41]). *A γ-certified algorithm for MIS satisfies the following:*
1. *returns the unique optimal solution, when run on a γ-stable instance,*
2. *is a γ-certified algorithm for Vertex Cover, and vice versa,*
3. *is a γ-approximation algorithm for MIS (and Vertex Cover).*

We stress that not all algorithms for stable instances are certified, so there is no equivalence between the two notions. Some examples (communicated to us by Yury Makarychev [43]) include the algorithms for stable instances of TSP [45], Max Cut (the GW SDP with triangle inequalities), and clustering. All these algorithms solve stable instances but are not certified. Thus, designing a certified algorithm is, potentially, a harder task than designing an algorithm for stable instances.

From now on, if an algorithm for MIS only returns a feasible solution S, it will be assumed to be "candidate" γ-certified that also returns the perturbed weight function w' with $w'_u = \gamma \cdot w_u$ for $u \in S$ and $w'_u = w_u$, otherwise.

3 Stable instances of MIS on special classes of graphs

In the next few sections, we obtain algorithms for stable instances of MIS on several natural classes of graphs, by using convex relaxations and combinatorial techniques.

3.1 Convex relaxations and robust algorithms

The starting point for the design of robust algorithms via convex relaxations is the structural result of Makarychev et al. [42], that gives sufficient conditions for the integrality of convex relaxations on stable instances. We now introduce a definition and restate their theorem in the setting of MIS.

▶ **Definition 7** ((α, β)-rounding). *Let $x : V \to [0, 1]$ be a feasible fractional solution of a convex relaxation of MIS whose objective value for an instance $G = (V, E, w)$ is $\sum_{u \in V} w_u x_u$. A randomized rounding scheme for x is an (α, β)-rounding, for some parameters $\alpha, \beta \geq 1$, if it always returns a feasible independent set S, such that the following two properties hold for every vertex $u \in V$:*
1. $\Pr[u \in S] \geq \frac{1}{\alpha} \cdot x_u$,
2. $\Pr[u \notin S] \leq \beta \cdot (1 - x_u)$.

▶ **Theorem 8** ([42]). *Let $x : V \to [0, 1]$ be an optimal fractional solution of a convex relaxation of MIS whose objective value for an instance $G = (V, E, w)$ is $\sum_{u \in V} w_u x_u$. Suppose that there exists an (α, β)-rounding for x, for some $\alpha, \beta \geq 1$. Then, x is integral for $(\alpha\beta)$-stable instances.*

The theorem suggests a simple robust algorithm: solve the relaxation, and if the solution is integral, report it, otherwise report that the instance is not stable (observe that the rounding scheme is used only in the analysis).

3.2 A robust algorithm for $(k - 1)$-stable instances of MIS on k-colorable graphs

In this section, we give a robust algorithm for $(k - 1)$-stable instances of MIS on k-colorable graphs by utilizing Theorem 8 and the standard LP for MIS. For a graph $G = (V, E, w)$, the standard LP has an indicator variable x_u for each vertex $u \in V$, and is given below.

$$\textbf{(LP)} \quad \max : \sum_{u \in V} w_u x_u \quad \text{s.t.:} \quad x_u + x_v \leq 1 \quad \forall (u, v) \in E, \quad \text{and} \quad x_u \in [0, 1] \quad \forall u \in V.$$

The corresponding polytope is half-integral [46], and so we always have an optimal solution x with $x_u \in \left\{0, \frac{1}{2}, 1\right\}$ for every $u \in V$. This is useful for designing (α, β)-rounding schemes, as it allows us to consider randomized combinatorial algorithms and easily present them as rounding schemes.

The crucial observation that we make is that the rounding scheme in Theorem 8 is only used in the analysis and is not part of the algorithm, and so it can run in super-polynomial time. We also note that the final (polynomial-time) algorithm does not need to have a k-coloring of the graph. Let $G = (V, E, w)$ be a k-colorable graph, and let x be an optimal half-integral solution. Let $V_i = \{u \in V : x_u = i\}$ for $i \in \{0, 1/2, 1\}$. We consider the rounding scheme of Hochbaum [35] (see Algorithm 1). We use the notation $[k] = \{1, ..., k\}$.

▶ **Theorem 9.** *Let $G = (V, E, w)$ be a k-colorable graph. Given an optimal half-integral solution x, the rounding scheme of Algorithm 1 is a $\left(\frac{k}{2}, \frac{2(k-1)}{k}\right)$-rounding for x.*

Algorithm 1 Hochbaum's k-colorable rounding scheme

1. Compute a k-coloring $f : V_{1/2} \to [k]$ of the induced graph $G[V_{1/2}]$.
2. Pick j uniformly at random from the set $[k]$, and set $V_{1/2}^{(j)} := \{u \in V_{1/2} : f(u) = j\}$.
3. Return $S := V_{1/2}^{(j)} \cup V_1$.

Proof. The set S is feasible, as there is no edge between V_1 and $V_{1/2}$ and f is a valid coloring. For $u \in V_0$, we have $\Pr[u \in S] = 0 = x_u$ and $\Pr[u \notin S] = 1 = 1 - x_u$. For $u \in V_1$, we have $\Pr[u \in S] = 1 = x_u$ and $\Pr[u \notin S] = 0 = 1 - x_u$. Let $u \in V_{1/2}$. We have $\Pr[u \in S] \geq \frac{1}{k} = \frac{2}{k} \cdot x_u$ and $\Pr[u \notin S] \leq 1 - \frac{1}{k} = \frac{2(k-1)}{k} \cdot (1 - x_u)$. The result follows. ◀

Theorems 8 and 9 now imply the following theorem, which is tight.

▶ **Theorem 10.** *The standard LP for* MIS *is integral for $(k-1)$-stable instances of k-colorable graphs.*

3.3 Algorithms for stable instances of MIS on bounded-degree graphs

Throughout this section, we assume that all graphs have maximum degree Δ. The only result (prior to our work) for stable instances on such graphs was using the greedy algorithm and was given by Bilu [15].

Algorithm 2 The greedy algorithm for MIS

1. Let $S := \emptyset$ and $X := V$.
2. while $(X \neq \emptyset)$:
   Set $S := S \cup \{u\}$ and $X := X \setminus (\{u\} \cup N(u))$, where $u := \arg\max_{v \in X}\{w_v\}$.
3. Return S.

▶ **Theorem 11** ([15])**.** *The greedy algorithm (see Algorithm 2) solves Δ-stable instances of* MIS *on graphs of maximum degree Δ.*

We first note that, since the maximum degree is Δ, the chromatic number is at most $\Delta + 1$, and so Theorem 10 implies a robust algorithm for Δ-stable instances, giving a robust analog of Bilu's result. In fact, we can slightly improve upon that by using Brook's Theorem [20], which states that the chromatic number is at most Δ, unless the graph is complete or an odd cycle. We can then prove following theorem.

▶ **Theorem 12.** *There exists a robust algorithm for $(\Delta - 1)$-stable instances of* MIS*, where Δ is the maximum degree.*

We now turn to non-robust algorithms and present an algorithm that solves $o(\Delta)$-stable instances, as long as the weights are polynomially-bounded integers. The core of the algorithm is a procedure that uses an α-approximation algorithm as a black-box in order to recover the optimal solution, when the instance is stable. Let $G = (V, E, w)$ be a graph with $n = |V|$ and $w : V \to \{1, ..., \texttt{poly}(n)\}$. Let $\mathcal{A}$ denote an α-approximation algorithm for MIS. We will give an algorithm for γ-stable instances with $\gamma = \left\lceil \sqrt{2\Delta\alpha} \right\rceil$. Note that we can assume that $\alpha \leq \Delta$ and $\gamma \leq \Delta$. These assumptions hold for the rest of this section. Algorithm 3 is the main algorithm, and it uses Algorithm 4 as a subroutine.

To prove the algorithm's correctness, we need some lemmas (see the full version [3] for their proofs).

Algorithm 3 Algorithm for γ-stable instances, where $\gamma = \left\lceil \sqrt{2\Delta\alpha} \right\rceil$

`Bounded-Alg`$(G(V,E,w))$:
1. If $w(V) \le \gamma$, then return V.
2. Run α-approximation algorithm $\mathcal{A}$ on G to get an independent set I.
3. Let $S :=$ `PURIFY`(G, I, γ).
4. Let $S' :=$ `Bounded-Alg`$(G[V \setminus (S \cup N(S))])$.
5. Return $S \cup S'$.

Algorithm 4 The `PURIFY` procedure

`INPUT`: Graph $G = (V,E,w)$, independent set $I \subseteq V$ and factor $\gamma \ge 1$.

1. Create a bipartite unweighted graph $G_0 = (L \cup R, E_0)$, where L contains $\gamma \cdot w(u)$ copies of each $u \in I$ and R contains $w(v)$ copies of each $v \in V \setminus I$. The set E_0 is defined as follows: if (u,v) is an edge in G with $u \in I$ and $v \notin I$, then add edges from each copy of u in L to each copy of v in R.
2. Compute a maximum cardinality matching M of G_0.
3. Return the set of all vertices $u \in I$ that have at least one unmatched copy in L w.r.t. M.

▶ **Lemma 13.** *Let $G = (V,E,w)$ be γ-stable, with $w_u \ge 1$, for every $u \in V$. If $w(V) \le \gamma$, then $E = \emptyset$.*

The above lemma justifies Step 1 of Algorithm 3.

▶ **Lemma 14.** *Let $G = (V,E,w)$ be γ-stable, and let I^* be its maximum independent set. Then $w(I^*) > \frac{\gamma}{2\Delta} \cdot w(V)$.*

▶ **Lemma 15.** *Let $G = (V,E,w)$ be a γ-stable instance, let I^* be its maximum independent set and let I' be an α-approximate independent set. Then $I^* \cap I' \neq \emptyset$.*

We now analyze the `PURIFY` procedure (Algorithm 4).

▶ **Lemma 16.** *Let G be a γ-stable instance that is given as input to the `PURIFY` procedure (see Algorithm 4), along with an α-approximate independent set I, and let I^* be its maximum independent set. If $I \neq I^*$, then the set S returned by the procedure always satisfies the following two properties:*
1. *$S \neq \emptyset$,*
2. *$S \subseteq I^*$.*

Proof. We first prove Property (1). Let's assume that $S = \emptyset$. This means that all vertices in L are matched. By construction, this implies that $\gamma \cdot w(I) \le w(V \setminus I)$. Since I is an α-approximation, we have that $\gamma \cdot w(I) \ge \frac{\gamma}{\alpha} \cdot w(I^*) > \frac{\gamma \cdot \gamma}{2\Delta\alpha} w(V) \ge \frac{2\Delta\alpha}{2\Delta\alpha} w(V) = w(V)$, where the second inequality is due to Lemma 14. We conclude that $w(V \setminus I) > w(V)$, which is a contradiction. Thus, $S \neq \emptyset$.

We turn to Property (2). Let $A = I \setminus I^*$ and $B = I^* \setminus I$. Let $A_0 \subseteq L$ be the copies of the vertices of set A in G_0, and let $B_0 \subseteq R$ be the copies of the vertices of set B in G_0. We will show that for every $Z \subseteq A_0$, we have $|N(Z) \cap B_0| \ge |Z|$. To see this, let $Z \subseteq A_0$, and let $I(Z) \subseteq A$ be the distinct vertices of A whose copies (not necessarily all of them) are included in Z. Since the instance is γ-stable, this implies that the weight of the neighbors $F \subseteq B$ of $I(Z)$ in I^* is strictly larger than $\gamma \cdot w(I(Z))$. By construction, we have that $|Z| \le \gamma \cdot w(I(Z))$, and the number of vertices in G_0 corresponding to vertices of F is equal to $w(F)$. Moreover, all of these $w(F)$ vertices are connected with at least one vertex in Z, which means that $w(F) = |N(Z) \cap B_0|$. This implies that $|N(Z) \cap B_0| > |Z|$. Thus, Hall's condition is satisfied, and so there exists a perfect matching between the vertices of A_0 and (a subset of the vertices of) B_0.

We observe now that the neighbors of all vertices in B_0 are only vertices in A_0 and not in $L \setminus A_0$. This means that any maximum matching matches all vertices of A_0 (otherwise, we could increase the size of the matching by matching all vertices in A_0). Thus, $S \subseteq I \cap I^*$. ◀

Putting everything together, and by using the $\widetilde{O}(\Delta/\log\Delta)$-approximation algorithm of Halldórsson [31] or Halperin [33] as a black-box, it is easy to prove the following theorem.

▶ **Theorem 17.** *Algorithm 3 correctly solves* $\left\lceil\sqrt{2\Delta\alpha}\right\rceil$*-stable instances in polynomial time. In particular, there is an algorithm that solves* $\widetilde{O}(\Delta/\sqrt{\log\Delta})$*-stable instances.*

3.4 Robust algorithms for $(1+\varepsilon)$-stable instances of `MIS` on planar graphs

In this section, we design a robust algorithm for $(1+\varepsilon)$-stable instances of `MIS` on planar graphs. Theorem 10 already implies a robust algorithm for 3-stable instances of planar `MIS`, but we will use the Sherali-Adams hierarchy (denoted as SA from now on) to reduce this threshold down to $1+\varepsilon$, for any fixed $\varepsilon > 0$. In particular, we show that $O(1/\varepsilon)$ rounds of SA suffice to optimally solve $(1+\varepsilon)$-stable planar instances. We will not introduce the SA hierarchy formally, and we refer the reader to the many available surveys about it (see, e.g., [24]). The t-th level of SA for `MIS` has a variable Y_S for every subset $S \subseteq V$ of size at most $|S| \leq t+1$, whose intended value is $Y_S = \prod_{u \in S} x_u$, where x_u is the indicator of whether u belongs to the independent set. The relaxation has size $n^{O(t)}$, and thus can be solved in time $n^{O(t)}$.

Our starting point is the work of Magen and Moharrami [40], which gives a SA-based PTAS for `MIS` on planar graphs, inspired by Baker's technique [7]. In particular, [40] gives a rounding scheme for the $O(t)$-th round of SA that returns a $(1+O(1/t))$-approximation. In this section, we slightly modify and analyze their rounding scheme, and prove that it satisfies the conditions of Theorem 8. For that, we need a theorem of Bienstock and Ozbay [14]. For any subgraph H of a graph $G = (V, E)$, let $V(H)$ denote the set of vertices contained in H.

▶ **Theorem 18** ([14]). *Let* $t \geq 1$ *and* Y *be a feasible vector for the* t*-th level SA relaxation of the standard Independent Set LP for a graph* G*. Then, for any subgraph* H *of* G *of treewidth at most* t*, the vector* $(Y_{\{u\}})_{u \in V(H)}$ *is a convex combination of independent sets of* H*.*

The above theorem implies that the t-th level SA polytope is equal to the convex hull of all independent sets of the graph, when the graph has treewidth at most t.

The rounding scheme of Magen and Moharrami [40]. Let $G = (V, E, w)$ be a planar graph and $\{Y_S\}_{S \subseteq V : |S| \leq t+1}$ be an optimal t-th level solution of SA. We denote $Y_{\{u\}}$ as y_u, for any $u \in V$. We first fix a planar embedding of G. V can then be naturally partitioned into sets $V_0, V_1, ..., V_L$, for some $L \in \{0, ..., n-1\}$, where V_0 is the set of vertices in the boundary of the outerface, V_1 is the set of vertices in the boundary of the outerface after V_0 is removed, and so on. Note that for any edge $(u, v) \in E$, we have $u \in V_i$ and $v \in V_j$ with $|i - j| \leq 1$. We will assume that $L \geq 4$, since, otherwise, the graph is at most 4-outerplanar and the problem can then be solved optimally [7].

Following [7], we fix a parameter $k \in \{1, ..., L\}$, and for every $i \in \{0, ..., k-1\}$, we define $B(i) = \bigcup_{j \equiv i (\text{mod } k)} V_j$. We now pick an index $j \in \{0, ..., k-1\}$ uniformly at random. Let $G_0 = G[V_0 \cup V_1 ... \cup V_j]$, and for $i \geq 1$, $G_i = G[\bigcup_{q=(i-1)k+j}^{ik+j} V_q]$, where for a subset $X \subseteq V$, $G[X]$ is the induced subgraph on X. Observe that every edge and vertex of G appears in one or two of the subgraphs $\{G_i\}$, and every vertex $u \in V \setminus B(j)$ appears in exactly one G_i.

Magen and Moharrami observe that for every subgraph $G_i = (V(G_i), E(G_i))$, the set of vectors $\{Y_S\}_{S \subseteq V(G_i):|S| \leq t+1}$ is a feasible solution for the t-th level SA relaxation of the graph G_i. This is easy to see, as the LP associated with G_i is weaker than the LP associated with G (on all common variables), since G_i is a subgraph of G, and this extends to SA as well. We need one more observation: a k-outerplanar graph has treewidth at most $3k-1$ (see [19]). By construction, each G_i is a $(k+1)$-outerplanar graph. Thus, by setting $t = 3k+2$, Theorem 18 implies that the vector $\{y_u\}_{u \in V(G_i)}$ can be written as a convex combination of independent sets of G_i. Let p_i be the corresponding distribution of independent sets of G_i, implied by $\{y_u\}_{u \in V(G_i)}$.

We now consider the following rounding scheme. For each G_i, we (independently) sample an independent set S_i of G_i according to p_i. Each vertex $u \in V \setminus B(j)$ belongs to exactly one G_i and is included in the final independent set S if $u \in S_i$. A vertex $u \in B(j)$ might belong to two different graphs G_i, G_{i+1}, and is included in S only if $u \in S_i \cap S_{i+1}$. The algorithm then returns S.

Before analyzing the algorithm, we note that standard tree-decomposition based arguments show that the rounding is constructive (i.e. polynomial-time; this fact is not needed for the algorithm for stable instances of planar `MIS`, but will be used when designing certified algorithms).

▶ **Theorem 19.** *The above randomized rounding scheme always returns a feasible independent set S, such that for every vertex $u \in V$,*

1. $\Pr[u \in S] \geq \frac{k-1}{k} \cdot y_u + \frac{1}{k} \cdot y_u^2$,
2. $\Pr[u \notin S] \leq \left(1 + \frac{1}{k}\right) \cdot (1 - y_u)$.

Proof. It is easy to see that S is always a feasible independent set. We now compute the corresponding probabilities. Since the marginal probability of p_i on a vertex $u \in G_i$ is y_u, for any fixed j, for every vertex $u \in V \setminus B(j)$, we have $\Pr[u \in S] = y_u$, and for every vertex $u \in B(j)$, we have $\Pr[u \in S] \geq y_u^2$. Since j is picked uniformly at random, each vertex $u \in V$ belongs to $B(j)$ with probability exactly equal to $\frac{1}{k}$. Thus, we conclude that for every vertex $u \in V$, we have $\Pr[u \in S] \geq \frac{k-1}{k} \cdot y_u + \frac{1}{k} \cdot y_u^2$, and $\Pr[u \notin S] \leq 1 - \left(\frac{k-1}{k} \cdot y_u + \frac{1}{k} \cdot y_u^2\right) = 1 - y_u + \frac{y_u}{k} \cdot (1 - y_u) \leq \left(1 + \frac{1}{k}\right) \cdot (1 - y_u)$. ◀

The above theorem implies that the rounding scheme is a $\left(\frac{k}{k-1}, \frac{k+1}{k}\right)$-rounding. The following theorem now is a direct consequence of Theorems 8 and 19.

▶ **Theorem 20.** *For every $\varepsilon > 0$, the SA relaxation of $\left(3\left\lceil \frac{2}{\varepsilon} \right\rceil + 5\right) = O(1/\varepsilon)$ rounds is integral for $(1+\varepsilon)$-stable instances of `MIS` on planar graphs.*

4 Stable instances of `MIS` on general graphs

In this section, we study stable instances of general graphs. We present a strong lower bound on any algorithm (not necessarily robust) that solves $o(\sqrt{n})$-stable instances. We complement this lower bound with an algorithm that solves (εn)-stable instances in time $n^{O(1/\varepsilon)}$.

4.1 Computational hardness of stable instances of MIS

We show that for general graphs it is unlikely to obtain efficient algorithms for solving γ-stable instances for small values of γ. Our hardness reduction is based on the *planted clique* conjecture [29, 44], which states that finding $o(\sqrt{n})$ sized planted independent sets/cliques in the Erdős-Rényi graph $G\left(n, \frac{1}{2}\right)$ is computationally hard. Let $G\left(n, \frac{1}{2}, k\right)$ denote the distribution over graphs obtained by sampling a graph from $G\left(n, \frac{1}{2}\right)$ and then picking a uniformly random subset of k vertices and deleting all edges among them. The conjecture is formally stated below.

▶ **Conjecture 21.** *Let $0 < \varepsilon < \frac{1}{2}$ be a constant. Suppose that an algorithm $\mathcal{A}$ receives an input graph G that is either sampled from the ensemble $G\left(n, \frac{1}{2}\right)$ or $G\left(n, \frac{1}{2}, n^{\frac{1}{2}-\varepsilon}\right)$. Then, no $\mathcal{A}$ that runs in time polynomial in n can decide, with probability at least $\frac{4}{5}$, which ensemble G was sampled from.*

Our lower bound follows from the observation that planted random instances are stable up to high values of γ, and this suffices to imply our main result.

▶ **Theorem 22.** *Let $\varepsilon > 0$ be a constant and consider a random graph G on n vertices generated by first picking edges according to the Erdős-Rényi model $G(n, \frac{1}{2})$, followed by choosing a set I of vertices of size $n^{\frac{1}{2}-\varepsilon}$, uniformly at random, and deleting all edges inside I. Then, with probability $1 - o(1)$, the resulting instance is a $\Theta(n^{\frac{1}{2}-\varepsilon}/\log n)$-stable instance of MIS.*

Proof. Let $G = (V, E)$ be the resulting graph (we assume that all weights are set to 1). We start by stating two well-known properties of the graph G that hold with probability $1 - o(1)$ ([2]).

1. For each vertex $u \in V \setminus I$, we have $|N(u) \cap I| \geq \frac{1}{2} \cdot n^{\frac{1}{2}-\epsilon} (1 \pm o(1))$.
2. The size of the maximum independent set in the graph $G[V \setminus I]$ is at most $\lceil 2(1 \pm o(1)) \log n \rceil$.

Consider any other independent set $S \neq I$. By Property 1, we have that $|I \setminus S| \geq \frac{1}{2} n^{\frac{1}{2}-\varepsilon}(1 - o(1))$. By Property 2, we must have that $|S \setminus I| \leq 2(1 \pm o(1)) \log n$. Hence, $|S| < |I|$ and furthermore, $|I \setminus S| > \gamma \cdot |S \setminus I|$ for $\gamma = \frac{n^{\frac{1}{2}-\varepsilon}}{4 \log n}$. We conclude that the instance is $\left(\frac{n^{\frac{1}{2}-\varepsilon}}{4 \log n}\right)$-stable. ◀

4.2 An algorithm for (ϵn)-stable instances

In this section, we design an algorithm for (εn)-stable instances on graphs of n vertices, that runs in time $n^{O(1/\varepsilon)}$; thus, $\varepsilon > 0$ is assumed to be constant. Due to space constraints, we only give an informal description of the algorithm for the special case of $(n/2)$-stable instances; the algorithm then naturally generalizes to (n/k)-stable instances, for any integer $k \geq 2$. The base case (i.e., $k = 1$) uses the greedy algorithm.

We start by observing that either the chromatic number is at most $n/2$, in which case the LP is integral, or there are more than $n/2$ vertices of degree at least $n/2$. In the latter case, we check whether each of these high-degree vertices belongs to the optimal solution, by removing each such vertex and its neighborhood, one at a time, and using the algorithm recursively. Since their neighborhoods are large, we end up with graphs with at most $n/2$ vertices, which are still $(n/2)$-stable if the removed vertex belongs to the optimal solution; thus, the recursion succeeds on them (in particular, the greedy algorithms solves such instances). Finally, if none of these high-degree vertices belongs to the optimal solution, then we remove all of them, and end up again with a graph with at most $n/2$ vertices that is $(n/2)$-stable; the recursive call again solves that instance. By returning the best of all computed solutions, we are guaranteed to recover the optimal solution.

▶ **Theorem 23.** *There exists an algorithm that solves $(\frac{n}{k})$-stable instances of* MIS *on graphs of n vertices in time $n^{O(k)}$.*

5 Stability and integrality gaps of convex relaxations

In this section, we state a general theorem about the integrality gap of convex relaxations of maximization problems on stable instances. In particular, we show that, even if the conditions of Theorem 8 are not satisfied, the integrality gap still significantly decreases as stability increases.

▶ **Theorem 24.** *Consider a relaxation for* MIS *that assigns a value $x_u \in [0,1]$ to each vertex u of a graph $G = (V, E, w)$, and its objective function is $\sum_{u \in V} w_u x_u$. Let α be its integrality gap, for some $\alpha > 1$. Then, its integrality gap is at most $\min\{\alpha, 1 + \frac{1}{\beta-1}\}$ on $(\alpha\beta)$-stable instances, for any $\beta > 1$.*

The proof is somewhat similar in spirit to the lemmas and analysis used for Algorithm 3. We stress that the above result is inherently non-constructive. Nevertheless, it suggests estimation algorithms for stable instances of MIS, such as the following, which is is a direct consequence of Theorem 24 and the results of [31, 33].

▶ **Corollary 25.** *For any fixed $\varepsilon > 0$, the Lovasz θ-function SDP has integrality gap at most $1 + \varepsilon$ on $\widetilde{O}\left(\frac{1}{\varepsilon} \cdot \frac{\Delta}{\log \Delta}\right)$-stable instances of* MIS *of maximum degree Δ.*

We note that the theorem naturally extends to many other maximization graph problems, and is particularly interesting for relaxations that require super-constant stability for the recovery of the optimal solution (e.g., the Max Cut SDP has integrality gap $1 + \varepsilon$ for $(2/\varepsilon)$-stable instances although the integrality gap drops to exactly 1 for $\Omega(\sqrt{\log n} \cdot \log \log n)$-stable instances).

In general, such a theorem is not expected to hold for minimization problems, but, in our case, MIS gives rise to its complementary Minimum Vertex Cover problem, and it turns out that we can prove a very similar result for Minimum Vertex Cover as well.

6 Certified algorithms for MIS

In this section, we initiate the systematic study of certified algorithms for MIS, introduced by Makarychev and Makarychev [41].

6.1 Certified algorithms using convex relaxations

An important observation that [41] makes is that an approach very similar to the one used for the design of algorithms for weakly-stable instances [42] can be used to obtain certified algorithms. More formally, they prove the following theorem.

▶ **Theorem 26** ([41])**.** *Let $x : V \to [0,1]$ be an optimal fractional solution of a convex relaxation of* MIS *whose objective value for an instance $G = (V, E, w)$ is $\sum_{u \in V} w_u x_u$. Suppose that there exists a polynomial-time (α, β)-rounding for x. Then, there exists a polynomial-time $(\alpha\beta + \varepsilon)$-certified algorithm for* MIS *on instances with integer polynomially-bounded weights (for $\varepsilon \geq 1/\text{poly}(n) > 0$).*

We now combine Theorem 19 with Theorem 26 and obtain the following theorem.

▶ **Theorem 27.** *There exists a polynomial-time $(1+\varepsilon)$-certified algorithm for* `MIS` *on planar graphs with integer polynomially-bounded weights (for $\varepsilon \geq 1/\text{poly}(n) > 0$).*

6.2 Combinatorial certified algorithms

In this section, we study several combinatorial algorithms for `MIS` and prove that they are certified. The first result is about the greedy algorithm.

▶ **Theorem 28.** *The greedy algorithm (see Algorithm 2) is a Δ-certified algorithm for* `MIS` *on graphs of maximum degree Δ. More generally, the greedy algorithm is a Δ-certified algorithm for any instance of a Δ-extendible system.*

Moreover, we introduce a variant of the greedy algorithm for `MIS` that is a $\sqrt{\Delta^2 - \Delta + 1}$-certified algorithm; the improvement over the greedy is moderate for small values of Δ. Finally, we show that the algorithm of Berman and Fürer [13] is $\left(\frac{\Delta+1}{3} + \varepsilon\right)$-certified, when all weights are 1. We acknowledge that the restriction to unweighted graphs limits the scope of the algorithm, but we consider this as a first step towards obtaining $(c\Delta)$-certified algorithms, for $c < 1$.

▶ **Theorem 29.** *The Berman-Fürer algorithm ([13]) is a $\left(\frac{\Delta+1}{3} + \varepsilon\right)$-certified algorithm for* `MIS` *on graphs of maximum degree Δ, when all weights are equal to 1.*

Let $G = (V, E, w)$ be a graph of maximum degree Δ, $n = |V|$, where $w_u = 1$ for every $u \in V$. We say X is an improvement of I, if both I and $I \oplus X$ are independent sets, the subgraph induced by X is connected and $I \oplus X$ is larger than I. (The operator $\oplus$ denotes the symmetric difference.)

The algorithm starts with a feasible independent set I' and iteratively improves the solution by checking whether there exists an improvement X with size $|X| \leq \sigma$. If so, it replaces I by $I \oplus X$ and repeats. Otherwise, if no such improvement exists, it outputs the current independent set I. Assuming that Δ is a constant, the algorithm runs in polynomial time as long as $\sigma = O(\log n)$.

▶ **Lemma 30** ([13]). *If Δ is a constant and $\sigma = O(\log n)$, the algorithm runs in polynomial time.*

The main result can be presented as follows. Along with Definition 3, it implies Theorem 29.

▶ **Lemma 31.** *Let I be the independent set returned by the algorithm with $\sigma = 32k\Delta^{4k} \log n$ and let $S \neq I$ be any feasible independent set. Then, we have $|S \setminus I| \leq \left(\frac{\Delta+1}{3} + \varepsilon\right) \cdot |I \setminus S|$, where $\varepsilon = \frac{1}{3k}$.*

Proof. Let $\bar{S} = S \setminus I$ and $\bar{I} = I \setminus S$. First, we observe that every $u \in \bar{S}$ has at least one neighbor in $\bar{I}$, otherwise, we could improve I by adding a new vertex from $\bar{S}$. We now consider the set $T = \{u \in \bar{S} : |N(u) \cap I| = 1\} \subseteq \bar{S}$. In words, T is the set of elements in $\bar{S}$ that have exactly one neighbor in I. We also define $J = \{v \in \bar{I} : N(v) \cap T \neq \emptyset\}$ to be the set of elements of $\bar{I}$ that have at least one neighbor in T. We will show that $|T| \leq |J|$.

To prove this, let's assume that $|T| > |J|$. Then, by the pigeonhole principle, we must have at least one vertex $v \in J$ that is connected to at least two vertices $u_1, u_2 \in T$. This implies that replacing v with u_1 and u_2 would be an improvement. Thus, we get a contradiction. Now let $I_0 = \bar{I} \setminus J$ and $S_0 = \bar{S} \setminus T$. The final step of the proof is a direct consequence of Lemma 3.5 of [13], that states that if there is no improvement over I of size at most $\sigma = 32k\Delta^{4k}\log n$, then for $\varepsilon = 1/(3k)$, $|S_0| \leq \left(\frac{\Delta+1}{3} + \varepsilon\right)|I_0|$. Recall that we have already proved $|T| \leq |J|$. Therefore,

$$|S \backslash I| = |S_0| + |T| \leq \left(\frac{\Delta+1}{3} + \varepsilon\right)|I_0| + |J| \leq \left(\frac{\Delta+1}{3} + \varepsilon\right)(|I_0| + |J|) = \left(\frac{\Delta+1}{3} + \varepsilon\right)|I \backslash S|.$$

◀

References

1 Noga Alon and Nabil Kahale. Approximating the independence number via the θ-function. *Math. Program.*, 80:253–264, 1998.

2 Noga Alon, Michael Krivelevich, and Benny Sudakov. Finding a large hidden clique in a random graph. *Random Structures and Algorithms*, 13(3-4):457–466, 1998.

3 Haris Angelidakis, Pranjal Awasthi, Avrim Blum, Vaggos Chatziafratis, and Chen Dan. Bilu-Linial stability, certified algorithms and the Independent Set problem. *CoRR*, abs/1810.08414, 2018. URL: `http://arxiv.org/abs/1810.08414`.

4 Haris Angelidakis, Konstantin Makarychev, and Yury Makarychev. Algorithms for stable and perturbation-resilient problems. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing (STOC)*, pages 438–451, 2017.

5 Per Austrin, Subhash Khot, and Muli Safra. Inapproximability of vertex cover and independent set in bounded degree graphs. *Theory of Computing*, 7(1):27–43, 2011.

6 Pranjal Awasthi, Avrim Blum, and Or Sheffet. Center-based clustering under perturbation stability. *Inf. Process. Lett.*, 112(1-2):49–54, 2012.

7 Brenda S. Baker. Approximation algorithms for NP-complete problems on planar graphs. *J. ACM*, 41(1):153–180, 1994.

8 Maria-Florina Balcan, Nika Haghtalab, and Colin White. k-center Clustering under Perturbation Resilience. In *Proceedings of the 43rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 68:1–68:14, 2016.

9 Maria-Florina Balcan and Yingyu Liang. Clustering under perturbation resilience. *SIAM J. Comput.*, 45(1):102–155, 2016.

10 Nikhil Bansal. Approximating independent sets in sparse graphs. In *Proceedings of the 26th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1–8, 2015.

11 Nikhil Bansal, Anupam Gupta, and Guru Guruganesh. On the Lovász Theta function for Independent Sets in sparse graphs. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, pages 193–200, 2015.

12 Nikhil Bansal, Daniel Reichman, and Seeun William Umboh. LP-based robust algorithms for noisy minor-free and bounded treewidth graphs. In *Proceedings of the 28th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1964–1979, 2017.

13 Piotr Berman and Martin Fürer. Approximating maximum independent set in bounded degree graphs. In *Proceedings of the 5th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 365–371, 1994.

14 Daniel Bienstock and Nuri Özbay. Tree-width and the Sherali-Adams operator. *Discrete Optimization*, 1(1):13–21, 2004.

15 Yonatan Bilu. On spectral properties of graphs and their application to clustering. *PhD Thesis*, pages 77–78, 2004.

16 Yonatan Bilu, Amit Daniely, Nati Linial, and Michael E. Saks. On the practically interesting instances of MAXCUT. In *Proceedings of the 30th International Symposium on Theoretical Aspects of Computer Science (STACS)*, pages 526–537, 2013.

17 Yonatan Bilu and Nathan Linial. Are stable instances easy? *Combinatorics, Probability & Computing*, 21(5):643–660, 2012.

18 Avrim Blum and Joel Spencer. Coloring random and semi-random k-colorable graphs. *Journal of Algorithms*, 19(2):204–234, 1995.

19 Hans L. Bodlaender. A partial k-arboretum of graphs with bounded treewidth. *Theoretical Computer Science*, 209(1):1–45, 1998.

20 R. L. Brooks. On colouring the nodes of a network. *Mathematical Proceedings of the Cambridge Philosophical Society*, 37(2):194–197, 1941.

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

22 Vaggos Chatziafratis, Tim Roughgarden, and Jan Vondrák. Stability and recovery for independence systems. In *Proceedings of the 25th Annual European Symposium on Algorithms (ESA)*, pages 26:1–26:15, 2017.

23 Chandra Chekuri and Shalmoli Gupta. Perturbation resilient clustering for k-center and related problems via LP relaxations. In *Proceedings of the 21st International Conference on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, 2018.

24 Eden Chlamtac and Madhur Tulsiani. *Convex Relaxations and Integrality Gaps*, pages 139–169. Springer US, Boston, MA, 2012.

25 Vincent Cohen-Addad and Chris Schwiegelshohn. On the local structure of stable clustering instances. In *Proceedings of the 58th IEEE Annual Symposium on Foundations of Computer Science (FOCS)*, pages 49–60, 2017.

26 Amit Deshpande, Anand Louis, and Apoorv Vikram Singh. On Euclidean k-means clustering with alpha-center proximity. In *the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 2087–2095, 2019.

27 Uriel Feige. Approximating maximum clique by removing subgraphs. *SIAM J. Discrete Math.*, 18(2):219–225, 2004.

28 Uriel Feige and Joe Kilian. Heuristics for semirandom graph problems. *Journal of Computer and System Sciences*, 63(4):639–671, 2001.

29 Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *J. ACM*, 64(2):8:1–8:37, 2017.

30 Zachary Friggstad, Kamyar Khodamoradi, and Mohammad R. Salavatipour. Exact algorithms and lower bounds for stable instances of Euclidean k-MEANS. In *Proceedings of the 30th ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 2958–2972, 2019.

31 Magnús M. Halldórsson. Approximations of weighted independent set and hereditary subset problems. *J. Graph Algorithms Appl.*, 4(1), 2000.

32 Magnús M. Halldórsson and Jaikumar Radhakrishnan. Improved approximations of independent sets in bounded-degree graphs via subgraph removal. *Nord. J. Comput.*, 1(4):475–492, 1994.

33 Eran Halperin. Improved approximation algorithms for the vertex cover problem in graphs and hypergraphs. *SIAM J. Comput.*, 31(5):1608–1623, 2002.

34 Johan Håstad. Clique is hard to approximate within $n^{1-\varepsilon}$. In *Proceedings of the 37th Annual Symposium on Foundations of Computer Science (FOCS)*, pages 627–636, 1996.

35 Dorit S. Hochbaum. Efficient bounds for the stable set, vertex cover and set packing problems. *Discrete Applied Mathematics*, 6(3):243 – 254, 1983.

36 Akihisa Kako, Takao Ono, Tomio Hirata, and Magnús M. Halldórsson. Approximation algorithms for the weighted independent set problem in sparse graphs. *Discrete Applied Mathematics*, 157(4):617–626, 2009.

37 Subhash Khot and Ashok Kumar Ponnuswami. Better inapproximability results for MaxClique, Chromatic Number and Min-3Lin-Deletion. In *Proceedings of the 33rd International Colloquium on Automata, Languages, and Programming (ICALP)*, pages 226–237, 2006.

38 Hunter Lang, David Sontag, and Aravindan Vijayaraghavan. Optimality of approximate inference algorithms on stable instances. In *the 21st International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 1157–1166, 2018.

39 Hunter Lang, David Sontag, and Aravindan Vijayaraghavan. Block stability for MAP inference. In *the 22nd International Conference on Artificial Intelligence and Statistics (AISTATS)*, pages 216–225, 2019.
40 Avner Magen and Mohammad Moharrami. Robust algorithms for Maximum Independent Set on minor-free graphs based on the Sherali-Adams hierarchy. In *Proceedings of the 12th International Workshop on Approximation Algorithms for Combinatorial Optimization Problems (APPROX)*, pages 258–271, 2009.
41 Konstantin Makarychev and Yury Makarychev. Certified approximation algorithms: Bridging worst-case and beyond-the-worst-case analysis. *Manuscript*, 2018.
42 Konstantin Makarychev, Yury Makarychev, and Aravindan Vijayaraghavan. Bilu-Linial stable instances of Max Cut and Minimum Multiway Cut. In *Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 890–906, 2014.
43 Yury Makarychev. Private communication. 2018.
44 Raghu Meka, Aaron Potechin, and Avi Wigderson. Sum-of-squares lower bounds for planted clique. In *Proceedings of the 47th Annual ACM on Symposium on Theory of Computing (STOC)*, pages 87–96, 2015.
45 Matús Mihalák, Marcel Schöngens, Rastislav Srámek, and Peter Widmayer. On the complexity of the metric TSP under stability considerations. In *SOFSEM 2011: Theory and Practice of Computer Science - 37th Conference on Current Trends in Theory and Practice of Computer Science*, pages 382–393, 2011.
46 G. L. Nemhauser and L. E. Trotter. Vertex packings: Structural properties and algorithms. *Mathematical Programming*, 8(1):232–248, 1975.
47 Leslie G. Valiant and Vijay V. Vazirani. NP is as easy as detecting unique solutions. *Theor. Comput. Sci.*, 47(3):85–93, 1986.
48 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. In *Proceedings of the 38th Annual ACM Symposium on Theory of Computing (STOC)*, pages 681–690, 2006.