

An algorithmic weakening of the Erdős-Hajnal conjecture

Édouard Bonnet

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Stéphan Thomassé

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Xuan Thang Tran

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Rémi Watrigant

Univ Lyon, CNRS, ENS de Lyon, Université Claude Bernard Lyon 1, LIP UMR5668, France

Abstract

We study the approximability of the MAXIMUM INDEPENDENT SET (MIS) problem in H -free graphs (that is, graphs which do not admit H as an induced subgraph). As one motivation we investigate the following conjecture: for every fixed graph H , there exists a constant $\delta > 0$ such that MIS can be $n^{1-\delta}$ -approximated in H -free graphs, where n denotes the number of vertices of the input graph. We first prove that a constructive version of the celebrated Erdős-Hajnal conjecture implies ours. We then prove that the set of graphs H satisfying our conjecture is closed under the so-called graph substitution. This, together with the known polynomial-time algorithms for MIS in H -free graphs (e.g. P_6 -free and fork-free graphs), implies that our conjecture holds for many graphs H for which the Erdős-Hajnal conjecture is still open. We then focus on improving the constant δ for some graph classes: we prove that the classical LOCAL SEARCH algorithm provides an $OPT^{1-\frac{1}{7}}$ -approximation in $K_{t,t}$ -free graphs (hence a \sqrt{OPT} -approximation in C_4 -free graphs), and, while there is a simple \sqrt{n} -approximation in triangle-free graphs, it cannot be improved to $n^{\frac{1}{4}-\varepsilon}$ for any $\varepsilon > 0$ unless $NP \subseteq BPP$. More generally, we show that there is a constant c such that MIS in graphs of girth γ cannot be $n^{\frac{c}{\gamma}}$ -approximated. Up to a constant factor in the exponent, this matches the ratio of a known approximation algorithm by Monien and Speckenmeyer, and by Murphy. To the best of our knowledge, this is the first strong (i.e., $\Omega(n^\delta)$ for some $\delta > 0$) inapproximability result for MAXIMUM INDEPENDENT SET in a proper hereditary class.

2012 ACM Subject Classification Theory of computation \rightarrow Graph algorithms analysis; Theory of computation \rightarrow Approximation algorithms analysis

Keywords and phrases Approximation, Maximum Independent Set, H -free Graphs, Erdős-Hajnal conjecture

1 Introduction

An *independent set* of a (simple, undirected) graph is a set of pairwise non-adjacent vertices. Independent sets have been central in various research topics, both in algorithmic and structural graph theory. In structural graph theory, independent sets (and their complements, cliques) are at the core of several celebrated results, such as König's theorem, Ramsey's theorem, or Turán's theorem [8], to name only a few. Finding an independent set of maximum cardinality (called the MAXIMUM INDEPENDENT SET problem, or MIS for short) is a fundamental intractable optimization problem. Indeed, it is NP-hard to solve [21], but also to approximate within ratio $n^{1-\varepsilon}$ for any $\varepsilon > 0$ [27, 35], where n denotes the number of vertices of the input graph. On the positive side, MIS becomes tractable when restricted to some specific graph classes: It is polynomial-time solvable in bipartite graphs and more generally in perfect graphs [22], admits a PTAS in planar graphs [6] and in more general geometric graph classes such as pseudo-disk graphs [10], bounded genus or H -minor-free

graphs [14]. Notice that all the aforementioned graph classes are closed under taking induced subgraphs. We call *hereditary* such a class, and add the qualificative *proper* if it is not the class of all graphs. A hereditary class can be defined by a (possibly infinite) set of forbidden induced subgraphs. A potentially unifying framework is to consider the complexity of MIS in H -free graphs (i.e., graphs without induced copy of H) and \mathcal{H} -free graphs (i.e., graphs without induced copy of any $H \in \mathcal{H}$). However, a classical reduction [2, 3] consisting of subdividing every edge of a given graph G a fixed even number of times $2c$ leads to a graph G' such that $\alpha(G') = \alpha(G) + c|E(G)|$ (where $\alpha(\cdot)$ denotes the size of a maximum independent set of a graph). This reduction, together with the fact that MIS remains APX-hard in graphs of maximum degree at most 3 [5] (which means in particular that we may assume that $\alpha(G) = \Omega(|E(G)|)$ in the reduction) implies the following:

► **Theorem 1** ([2, 3] and [5]). *For any fixed connected graph H which is neither a path nor a subdivision of the claw $K_{1,3}$, MIS is APX-hard in H -free graphs.*

On the positive side, polynomial algorithms are known for P_6 -free graphs [24] and fork-free graphs [4]. For paths on at least seven vertices and subdivided claws not contained in the fork, the computational complexity of MIS remains unsettled.

In this work, we start a systematic investigation of the approximability of MIS in H -free graphs. The intuition is that forbidding a fixed graph H as an induced subgraph should imply a drastic change in the structure of independent sets and cliques. This idea is at the core of the Erdős-Hajnal conjecture: while in random graphs of $\mathcal{G}(n, 1/2)$ the expected maximum of the clique number and the independence number is $O(\log n)$ [19], this value should be significantly larger for an H -free graph. More formally:

► **Definition 2.** *A graph H satisfies the Erdős-Hajnal property if there exists a constant $\delta > 0$ such that every H -free graph G with n vertices contains either a clique or an independent set of size n^δ .*

► **Conjecture 3** ([18]). *Every graph H satisfies the Erdős-Hajnal property.*

So far, the Erdős-Hajnal conjecture has been verified for only a small number of graphs, namely: all graphs on at most four vertices, the bull¹, the cliques, and every graph that can be constructed from them using the so-called *substitution* operation [12] (we describe this operation in Section 2). Interestingly, for many graphs H satisfying the Erdős-Hajnal property, MIS is known to be either polynomial or at least to admit an $n^{1-\varepsilon}$ -approximation algorithm for some $\varepsilon > 0$. A typical example of this situation is when H is the clique of size $t > 1$. In that case, Ramsey's theorem can be invoked to get a $n^{\frac{t-2}{t-1}}$ -approximation algorithm. Indeed a K_t -free graph always contains an independent set of size at least $n^{\frac{1}{t-1}}$, and the classical proof readily yields a polytime algorithm finding such an independent set. This leads us to define an approximation weaker version of the Erdős-Hajnal property and its companion conjecture:

► **Definition 4.** *A graph H satisfies the improved approximation property if there exists a constant $\varepsilon > 0$ such that MIS admits a (randomized) $n^{1-\varepsilon}$ -approximation polynomial algorithm on every H -free n -vertex graph G .*

Here, a *randomized* ρ -approximation algorithm is an algorithm which, given an input graph on n vertices, outputs a ρ -approximation of the problem *with high probability* (*w.h.p.*

¹ The bull is the graph obtained by adding a pending vertex to two different vertices of a triangle.

for short), that is with probability at least a function of n tending to 1 when n goes to infinity.

► **Conjecture 5.** *Every graph H satisfies the improved approximation property.*

We refer to Conjecture 5 as the *improved approximation conjecture*. Informally, it states that the inapproximability of MIS in general graphs can be beaten in any proper hereditary class.

Results and organization of the paper.

On the one hand, there exist graphs H satisfying the improved approximation property for which the Erdős-Hajnal conjecture is still open. Indeed, as mentioned previously, MIS is polynomial-time solvable in P_6 -free graphs, whereas it is still open whether P_5 satisfies the Erdős-Hajnal property. On the other hand, one may wonder if the satisfiability of the Erdős-Hajnal property for a graph H can help designing an approximation algorithm in H -free graphs, and more concretely if Conjecture 3 implies Conjecture 5. In Section 2, we prove that this is almost the case. More precisely, we prove that every graph H satisfying a *constructive* version of the Erdős-Hajnal property also satisfies the improved approximation property. We also show that the improved approximation property is preserved through the substitution operation, which is the one graph operation known to preserve the Erdős-Hajnal property.

We then try and obtain better approximation ratios for the improved approximation property: for a given H , what is the largest $\varepsilon > 0$ such that MIS admits an $O(n^{1-\varepsilon})$ -approximation algorithm in H -free graphs? We investigate this question in Sections 3 and 4. More precisely, in Section 3 we describe some particular properties of graphs H as well as graph operations preserving the improved approximation property in a better way than the substitution. We also prove that the classical local search algorithm provides a \sqrt{OPT} approximation ratio in C_4 -free graphs and, more generally, an $O(OPT^{1-1/t})$ -approximation algorithm in $K_{t,t}$ -free graphs. Finally, we present in Section 4 some negative results concerning the improved approximation property: while MIS can be easily $n^{1/2}$ -approximated in triangle-free graphs, we show that this ratio cannot be improved to $n^{1/4-\varepsilon}$ for any $\varepsilon > 0$, unless $NP \subseteq BPP$. We also provide a generalization of this result when we forbid all cycles of length $3, \dots, t$ for a fixed $t \geq 3$.

Notations and definitions.

For two positive integers $i < j$, we denote the set of integers at least i and at most j by $[i, j]$, while $[i]$ is a short-hand for $[1, i]$. All the graphs we consider are simple; they have no multiple edges nor loops. For a vertex v in a simple graph G , $N_G(v)$, or simply $N(v)$ if the graph is unambiguous, denotes the set of neighbors of v . The *closed neighborhood* of v is defined as $N[v] := N(v) \cup \{v\}$. A *universal vertex* is a vertex whose closed neighborhood is the entire set of vertices. The size of a maximum independent set of G is denoted by $\alpha(G)$. The *girth* (resp. *odd girth*) of a graph is the smallest size of an induced cycle (resp. odd cycle) in the graph. K_s , P_s , C_s respectively denotes the clique, the path, and the cycle on s vertices, and $K_{s,t}$ is the biclique with s vertices on one side and t on the other side. The graph $K_3 = C_3$ is also called the *triangle*. The *claw* is the biclique $K_{1,3}$. The *fork* is the 5-vertex graph obtained by subdividing one edge of the claw. For a triple of integers $0 \leq i \leq j \leq k$, the graph $S_{i,j,k}$ is obtained by subdividing one edge of a claw $i - 1$ times, a second edge, $j - 1$ times, and a third edge $k - 1$ times (with the convention that subdividing

–1 times means removing the edge and its degree-one endpoint). Observe that with that definition, the family $\{S_{i,j,k}\}_{0 \leq i \leq j \leq k}$ contains the paths.

2 Constructive Erdős-Hajnal and the substitution operation

A graph H is said to satisfy the *constructive Erdős-Hajnal* property if there is a constant $\delta > 0$ and a polynomial-time algorithm which takes as input an H -free graph G , and outputs a clique or an independent set of size at least $|V(G)|^\delta$. We prove that the constructive Erdős-Hajnal conjecture implies Conjecture 5. To our knowledge, all the graphs H shown to satisfy the Erdős-Hajnal property so far, also satisfy its constructive version.

► **Theorem 6.** *Let H be a graph which satisfies the constructive Erdős-Hajnal property with constant² $0 < \delta \leq 1/2$. Then H satisfies the improved approximation property with constant $\delta - \delta^2 - \varepsilon$ for any fixed $\varepsilon > 0$.*

Proof. Let G be an H -free graph with $n := |V(G)|$. We assume $n \geq 2^{\frac{1}{1-(\delta-\delta^2)}}$, since otherwise the problem can be solved optimally in constant time. We prove that the algorithm described in Figure 1 provides a $n^{1-(\delta-\delta^2)}$ -approximation. In this algorithm, **Constructive-Erdős-Hajnal**(J) represents the polynomial-time algorithm which takes a graph J and outputs a set of at least $|V(J)|^\delta$ vertices of J which is either an independent set or a clique.

Input: a graph G

Output: an independent set of G

```

1:  $V' \leftarrow V(G)$ 
2: while  $|V'| \geq n^{1-\delta}$  do
3:    $X \leftarrow \text{Constructive-Erdős-Hajnal}(G[V'])$ 
4:   if  $X$  is an independent set of  $G$  then
5:     return  $X$ 
6:   else
7:      $V' \leftarrow V' \setminus X$ 
8: return  $\{v\}$ , for an arbitrary chosen  $v \in V(G)$ 

```

■ **Figure 1** Approximation algorithm for MIS in H -free graphs satisfying the constructive Erdős-Hajnal property.

Let X be the independent set returned by the algorithm. If X is returned through line 5, then by the definition of the **Constructive-Erdős-Hajnal** algorithm, we have $|X| \geq n^{(1-\delta)\delta}$ which is obviously an $n^{1-(\delta-\delta^2)}$ -approximate solution, since any optimal solution has size at most n .

Otherwise, X is returned through line 8 and is thus of size 1. However, in this case, observe that $V(G)$ is partitioned into cliques C_1, \dots, C_q , and the last set V' . Observe that $|V'| < n^{1-\delta}$, and that $|C_i| \geq n^{\delta-\delta^2}$ for every $i \in \{1, \dots, q\}$. We thus have $q \leq n^{1-(\delta-\delta^2)}$.

² Notice that MIS in H -free graphs is trivial if H has at most 2 vertices, whereas any graph with at least three vertices cannot satisfy the Erdős-Hajnal property with a constant $\delta > 1/2$. This is the reason why we assume $0 < \delta \leq 1/2$.

But also observe that in that case:

$$\begin{aligned}
 \alpha(G) &\leq q + |V'| \\
 &\leq n^{1-(\delta-\delta^2)} + n^{1-\delta} \\
 &\leq 2n^{1-(\delta-\delta^2)} \quad \text{since } n \geq 2^{\frac{1}{1-(\delta-\delta^2)}} \\
 &\leq n^{1-(\delta-\delta^2)+\varepsilon} \quad \text{as we may assume } n \geq 2^{1/\varepsilon} \text{ for any fixed } \varepsilon > 0.
 \end{aligned}$$

◀

It is natural to ask which kind of graph operations preserves the satisfiability of the improved approximation property. Given the previous result, natural candidates are graph operations preserving the Erdős-Hajnal property. In the following we prove that this is indeed the case concerning the substitution operation.

► **Definition 7.** Let H_1, H_2 be two vertex-disjoint graphs and $v_0 \in V(H_1)$. We say that a graph H is obtained from H_1 by substituting H_2 at v_0 if:

- $V(H) = (V(H_1) \setminus \{v_0\}) \cup V(H_2)$
- For $v, v' \in V(H_1) \setminus \{v_0\}$, vv' is an edge in H if and only if it is an edge in H_1 .
- For $v, v' \in V(H_2)$, vv' is an edge in H if and only if it is an edge in H_2 .
- For $v \in V(H_1) \setminus \{v_0\}$, $v' \in V(H_2)$, vv' is an edge in H if and only if vv_0 is an edge in H_1 .

More generally, we say that a graph H is obtained from H_1 and H_2 by substitution if there exists $v_0 \in V(H_1)$ such that H is obtained from H_1 by substituting H_2 at v_0 .

► **Theorem 8.** Let H_1, H_2 be two fixed graphs satisfying the improved approximation property. Then every graph H obtained from H_1 and H_2 by substitution satisfies the improved approximation property.

Let us start by sketching the idea of our algorithm. We first check whether the number of copies of H_1 in G is small. If so, then a randomly chosen subset of vertices of appropriate size will be H_1 -free *w.h.p.*, and we will be able to run our approximation algorithm for H_1 -free graphs. If the number of copies of H_1 is large, then we claim that we can find a large subset of vertices inducing an H_2 -free graph, and we thus run our approximation algorithm for H_2 -free graphs. Each time we run one of our approximation algorithms in an induced subgraph $G[X]$ which is either H_1 -free or H_2 -free, either it outputs a solution of size at least n^δ for some constant δ , in which case we are done, or it means that $\alpha(G[X])$ is small, in which case we keep X apart and continue the algorithm on $G[V \setminus X]$ as long as enough vertices survive. If too many vertices were kept apart along the process, it means that $\alpha(G)$ was very small at the beginning, so that any singleton $\{v\}$ is actually an approximated solution. We now prove formally the result.

Proof. Let $approx_{H_1}(G)$ (resp. $approx_{H_2}(G)$) be a polynomial-time algorithm which takes as input an H_1 -free graph (resp. H_2 -free graph) G on n vertices and outputs an $n^{1-\varepsilon_1}$ (resp. $n^{1-\varepsilon_2}$)-approximated solution for the MIS problem in G , for some $\varepsilon_1 > 0$ (resp. $\varepsilon_2 > 0$). For the sake of readability, we set $\varepsilon = \min\{\varepsilon_1, \varepsilon_2, 0.99\}$, so that $approx_{H_1}$ and $approx_{H_2}$ are $n^{1-\varepsilon}$ -approximation algorithms in H_1 -free graphs and H_2 -free graphs, respectively³.

³ Our result also holds if $approx_{H_1}$ and $approx_{H_2}$ are exact algorithms (hence with $\varepsilon_1 = \varepsilon_2 = 1$), but, for technical reasons, we view them as $n^{0.01}$ -approximation algorithms.

Let H be the graph obtained by substituting H_2 at some vertex $v_0 \in V(H_1)$, and let us consider an H -free graph G . We denote by n , n_1 and n_2 the number of vertices of G , H_1 and H_2 , respectively. We say that $X \subseteq V(G)$ is a set of H_1 -candidates if there exists a set $K \subseteq V(G)$ of $n_1 - 1$ vertices such that $G[K]$ is isomorphic to $H_1 - \{v_0\}$ and, for every $x \in X$, $G[K \cup \{x\}]$ is isomorphic to H_1 . Since G is H -free, $G[X]$ is H_2 -free.

Let $\gamma = \frac{\varepsilon}{2n_1}$, $\eta = \min(1 - \varepsilon, \gamma)$, and $\delta = \frac{\varepsilon\eta}{2 + \varepsilon\eta}$. We prove that the algorithm described in Figure 2 is an $O(n^{1-\delta})$ -approximation algorithm for MIS in H -free graphs.

Input: an H -free graph G with n vertices

Output: an independent set of G

```

1:  $i = 1, V_1 \leftarrow V(G)$ 
2: while  $|V_i| \geq n^{1-\delta}$  do
3:   if  $G[V_i]$  contains less than  $|V_i|^{n_1-\varepsilon}$  copies of  $H_1$  then
4:     pick a set  $X_i \subseteq V_i$  of size  $\lceil |V_i|^\gamma \rceil$  uniformly at random
5:     if  $G[X_i]$  is  $H_1$ -free then  $\triangleright$  This condition is true w.h.p.
6:        $W_i \leftarrow \text{approx}_{H_1}(G[X_i])$ 
7:       if  $|W_i| \geq n^\delta$  then return  $W_i$ 
8:     else return FAIL
9:   else
10:    find a set of  $H_1$ -candidates  $X_i \subseteq V_i$  with  $|X_i| \geq |V_i|^{1-\varepsilon}$ 
11:     $W_i \leftarrow \text{approx}_{H_2}(G[X_i])$   $\triangleright G[X_i]$  is  $H_2$ -free
12:    if  $|W_i| \geq n^\delta$  then return  $W_i$ 
13:     $V_{i+1} \leftarrow V_i \setminus X_i$ 
14:     $i \leftarrow i + 1$ 
15: return  $\{v\}$  for an arbitrary vertex  $v \in V$ 

```

■ **Figure 2** Approximation algorithm for MIS in H -free graphs, where H is the substitution of H_1 and H_2 .

► **Lemma 9.** *Algorithm 2 runs in polynomial time.*

Proof. An important remark is that at every step i , the graph $G[V_i]$ is an induced subgraph of G , hence is H -free. In line 3 (resp. 5), the algorithm runs through all subsets of n_1 vertices of V_i (resp. X_i), which can be done in $O(n^{n_1})$ time.

Finally, the existence of a set of H_1 -candidates in line 10 is ensured by the fact that in that case, $G[V_i]$ contains at least $|V_i|^{n_1-\varepsilon}$ copies of H_1 . Hence, by the pigeonhole principle, there must exist $n_1 - 1$ vertices $V_H \subseteq V_i$ such that $G[V_H]$ induces $H_1 - v_0$ together with a set $X_i \subseteq V_i \setminus V_H$ of size at least $|V_i|^{1-\varepsilon}$ such that for every $x \in X_i$, $G[V_H \cup \{x\}]$ induces H_1 . Finding the set V_H can be done in $O(|V_i|^{n_1-1})$ time, while finding the set X_i can be done in $O(|V_i|)$ time, since it is sufficient to find the vertices in $V_i \setminus V_H$ with the right neighborhood with respect to V_H . By the definition of H_1 -candidates, $G[X_i]$ is H_2 -free, which allows to run approx_{H_2} on $G[X]$ in the next line of the algorithm. ◀

We now prove that, *w.h.p.*, the solution S returned by our algorithm is an $O(n^{1-\delta})$ -approximation. To this end, we first prove that *w.h.p.* it does not return FAIL.

► **Lemma 10.** *If the number of copies of H_1 in a graph G on n vertices is less than $n^{n_1-\varepsilon}$, then any subset of vertices of size $\lceil n^\gamma \rceil$ picked uniformly at random induces an H_1 -free graph, with high probability.*

Proof. Let n be the number of vertices of G , and P be a subset of $[n^\gamma]$ vertices picked uniformly at random. For any set $V_H \subseteq V$ inducing H_1 , the probability that V_H is contained in P is $\frac{\binom{n-n_1}{|P|-n_1}}{\binom{n}{|P|}} < \left(\frac{|P|}{n}\right)^{n_1}$. Hence the probability that P is H_1 -free is at least

$$\left(1 - \left(\frac{|P|}{n}\right)^{n_1}\right)^{n^{n_1-\varepsilon}} = \left(1 - \frac{1}{n^{n_1-\frac{\varepsilon}{2}}}\right)^{n^{n_1-\varepsilon}}$$

which tends to 1 when $n \rightarrow +\infty$. ◀

Next, if it returns a solution through lines 7 or 12, then this solution is an independent set of size at least n^δ , by definition. We now deal with the case in which it returns a singleton, through line 15. The aim is to prove that $\alpha(G)$ is at most $O(n^{1-\delta})$. Let $q+1$ be the largest value of i in the execution of the algorithm (i.e., $|V_{q+1}| < n^{1-\delta}$). The vertex-set V is thus partitioned into X_1, \dots, X_q , and V_{q+1} . Hence we have $\alpha(G) \leq |V_{q+1}| + \sum_{i=1}^q \alpha(G[X_i])$. Since $|V_{q+1}| < n^{1-\delta}$, we only need to upper bound the second part.

► **Lemma 11.** *With the above definitions, $\sum_{i=1}^q \alpha(G[X_i]) \leq n^{1-\delta}$.*

Proof. Recall that we have $X_i \subseteq V_i$, where $V_i = V \setminus \bigcup_{j=1}^{i-1} X_j$, and, for each $i \in [q]$, we have constructed an independent set $W_i \subseteq X_i$. All these sets have the following properties:

1. $|V_{q+1}| < n^{1-\delta}$, by definition of q .
2. $|X_i| \geq |V_i|^\eta \geq n^{\eta(1-\delta)}$. Indeed, if X_i is defined in line 4, then it is of size at least $|V_i|^\gamma$, whereas if it is defined in line 10, it is of size at least $|V_i|^{1-\varepsilon}$, and $\eta = \min(1-\varepsilon, \gamma)$.
3. $|W_i| < n^\delta$, otherwise we would have returned it.
4. $\alpha(G[X_i]) \leq |W_i| \cdot |X_i|^{1-\varepsilon}$, since W_i is returned by *approx_{H1}* or *approx_{H2}*, which are approximation algorithms applied to $G[X_i]$.

Now, we have the following:

$$\begin{aligned} \sum_{i=1}^q \alpha(G[X_i]) &\leq \sum_{i=1}^q |W_i| \cdot |X_i|^{1-\varepsilon} \\ &\leq n^\delta \sum_{i=1}^q |X_i|^{1-\varepsilon} \end{aligned}$$

We then need the following technical lemma.

► **Lemma 12.** *Let $(a_i)_{i=1\dots q}$ be some positive numbers (with $q \in \mathbb{N}$) such that $\sum_{i=1}^q a_i = N$ and $a_i \geq k > 0$ for all $i \in \{1, \dots, q\}$. Then $\sum_{i=1}^q a_i^\zeta \leq Nk^{\zeta-1}$ for any $0 < \zeta < 1$.*

Proof. We have:

$$\sum_{i=1}^q a_i^\zeta \leq \left(\frac{N}{q}\right)^\zeta \cdot q = N \cdot \left(\frac{N}{q}\right)^{\zeta-1} \leq Nk^{\zeta-1}.$$

Using the above lemma together with item 2 of the previous properties in order to lower bound each $|X_i|$, we obtain:

$$\begin{aligned}
\sum_{i=1}^q \alpha(G[X_i]) &\leq n^\delta \left(\sum_{i=1}^q |X_i| \right) n^{\eta(1-\delta)(1-\varepsilon-1)} \\
&\leq n^{1-\varepsilon\eta(1-\delta)+\delta} \quad \text{since } \sum_{i=1}^q |X_i| \leq n \\
&\leq n^{1-\delta} \quad \text{because } \delta = \frac{\varepsilon\eta}{2+\varepsilon\eta}, \text{ hence } \varepsilon\eta(1-\delta) = 2\delta
\end{aligned}$$

◀

Hence, any solution of size 1 is an $O(n^{1-\delta})$ -approximation in this case, which concludes the proof. ◀

3 Better approximation ratios

In this section we improve over the ratio given by Theorem 8 for some graphs H that can be built by a sequence of substitutions from graphs H' such that MIS is polynomial-time solvable in H' -free graphs. Furthermore, we present deterministic algorithms.

3.1 Adding a universal vertex


Let H^{+u} be the graph H augmented by a universal vertex, i.e., we add one vertex adjacent to all the vertices of H .

► **Lemma 13.** *Let $0 \leq \gamma < 1$ be a real number and H be a graph such that MIS admits an OPT^γ -approximation \mathcal{A} in H -free graphs. Then it also admits an $\text{OPT}^{\frac{1}{2-\gamma}}$ -approximation \mathcal{A}^{+u} in H^{+u} -free graphs.*

Proof. Let G be the input graph, thus $\text{OPT} := \alpha(G)$. The base case of the algorithm is when G does not contain any vertex, and we correctly report the empty set as optimum solution. Otherwise G has at least one vertex, say v_1 . We run the approximation \mathcal{A} on $G[N(v_1)]$. G being H^{+u} -free, the subgraph induced by the open neighborhood of any vertex is indeed H -free. Let S_1 be the returned solution. By assumption, $|S_1| \geq \alpha(G[N(v_1)])^{1-\gamma}$. For what follows, the knowledge of the value OPT would help. Unfortunately we will make some recursive calls to \mathcal{A}^{+u} , so exhaustively guessing this value would result in an exponential running time. Instead we will branch but the branching tree will only have at most $n := |V(G)|$ leaves. More precisely the tree will be a so-called *comb*, i.e., a path where all the vertices except one end has an additional private neighbor. We eventually output the best solution found among all the leaves.

We inductively run \mathcal{A}^{+u} on $G - N[v_1]$, which produces a tree T with at most $n - 1$ leaves. And we output the best solution among S and all the solutions at the leaves of T augmented by the vertex v_1 . This algorithm returns an independent set since v_1 is by definition non-adjacent to any vertex of $G - N[v_1]$. The running time of our algorithm satisfies $f_{\mathcal{A}^{+u}}(n) = f_{\mathcal{A}}(n_1 - 1) + f_{\mathcal{A}^{+u}}(n - n_1) + O(1)$ (with $n_1 = |N(v_1)|$). Hence $f_{\mathcal{A}^{+u}}(n) = O(\max\{f_{\mathcal{A}}(n), n\})$, and in the likely event that \mathcal{A} is *not* sublinear, \mathcal{A}^{+u} has the same running time as \mathcal{A} up to a multiplicative constant factor.

We shall now show that \mathcal{A}^{+u} is indeed a $\text{OPT}^{\frac{1}{2-\gamma}}$ -approximation. We denote by v_1, \dots, v_p with $p \leq n$, the vertices added along the path to the deepest leaf of T . We denote by S_1, \dots, S_p the sets returned by \mathcal{A} such that S_i is computed in the graph $G'_i := G[N(v_i) \setminus \bigcup_{j < i} N[v_j]]$.



■ **Figure 3** The three maximal locally easy graphs H constructed from P_6 , the fork, and $tK_{1,3}$, respectively.

We also define $G_i := G[N[v_i] \setminus \bigcup_{j < i} N[v_j]]$, and $R_i := G - \bigcup_{j < i} N[v_j]$. Observe that $\{V(G_1), \dots, V(G_p)\}$ is a partition of $V(G)$, as well as, $\{V(G_1), \dots, V(G_i), V(R_{i+1})\}$ for every $i \in [p-1]$.

Let, if it exists, S_h be the first solution returned by \mathcal{A} when called on G'_h such that $|S_h| \geq \alpha(R_h)^{1-\frac{1}{2-\gamma}}$. We claim that the solution output at this leaf, namely $S'_h := S_h \cup \{v_1, \dots, v_{h-1}\}$ is an $\text{OPT}^{\frac{1}{2-\gamma}}$ -approximation. If such an S_h does not exist, we show the same statement where $S_h = \emptyset$ and $h-1 = p$.

We upperbound $\alpha(G_i)$ for every $i \in [h-1]$. By definition of S_h , it holds that $|S_i| < \alpha(R_i)^{1-\frac{1}{2-\gamma}}$ for any $i \in [h-1]$. Due to the approximation ratio of \mathcal{A} , it holds that:

$$\alpha(G_i)^{1-\gamma} \leq |S_i| < \alpha(R_i)^{1-\frac{1}{2-\gamma}} = \alpha(R_i)^{\frac{1-\gamma}{2-\gamma}},$$

hence $\alpha(G_i) < \alpha(R_i)^{\frac{1}{2-\gamma}} \leq \alpha(G)^{\frac{1}{2-\gamma}}$. Thus,

$$\begin{aligned} \text{OPT} &= \alpha(G) \leq \alpha(R_h) + \sum_{i \in [h-1]} \alpha(G_i) \leq |S_h| \alpha(R_h)^{\frac{1}{2-\gamma}} + \sum_{i \in [h-1]} \alpha(G)^{\frac{1}{2-\gamma}} \\ &\leq |S_h| \alpha(G)^{\frac{1}{2-\gamma}} + (h-1) \alpha(G)^{\frac{1}{2-\gamma}} = (|S_h| + h-1) \alpha(G)^{\frac{1}{2-\gamma}} = |S'_h| \alpha(G)^{\frac{1}{2-\gamma}} = |S'_h| \text{OPT}^{\frac{1}{2-\gamma}}. \end{aligned}$$

Therefore \mathcal{A}^{+u} is an $\text{OPT}^{\frac{1}{2-\gamma}}$ -approximation for MIS in H^{+u} -free graphs. ◀

3.2 Locally easy graphs

We say that a graph H is *locally easy* if it has a universal vertex v such that there is a polynomial-time algorithm for MIS in $H - \{v\}$ -free graphs. Up to now, the three maximal graphs H for which we know that MIS is polynomial-time solvable on H -free graphs are P_6 [23], the fork [4, 30], and $tK_{1,3}$ (or *tclaw*) [9] (see Figure 3 for the corresponding maximal locally easy graphs).

The following is an immediate consequence of Lemma 13. Therein we recall that the constant γ may take value 0.

► **Theorem 14.** *For any locally easy H , MIS can be $\sqrt{\text{OPT}}$ -approximated on H -free graphs.*

And in particular, there is a \sqrt{n} -approximation for triangle-free graphs. Lemma 13 also yields the following approximation ratio in K_{t+1} -free graphs.

► **Theorem 15.** *For any $t \geq 1$, MIS can be $\text{OPT}^{1-\frac{1}{t}}$ -approximated on K_{t+1} -free graphs.*

Proof. We show this statement by induction. The base case says that we can exactly solve in polynomial time MIS in edgeless graphs, which is obviously true. We assume that the statement is true for a fixed t . As $K_{t+2} = (K_{t+1})^{+u}$, Lemma 13 implies that MIS can be $\text{OPT}^{\frac{1}{2-(1-\frac{1}{t})}}$ -approximated on K_{t+2} -free graphs. Furthermore, $\frac{1}{2-(1-\frac{1}{t})} = \frac{1}{1+\frac{1}{t}} = \frac{t}{t+1} = 1 - \frac{1}{t+1}$. Therefore we do obtain an $\text{OPT}^{1-\frac{1}{t+1}}$ -approximation on K_{t+2} -free graphs. ◀

We say that a graph H is t -locally easy if it has a set U of t universal vertices such that a polynomial algorithm is known for MIS in $H - U$ -free graphs. Informally, these graphs are obtained by replacing universal vertices of Figure 3 by a k -clique. The previous result readily generalizes from K_{t+1} -free to H -free graphs with H $(t + 1)$ -locally easy, with the same proof.

► **Corollary 16.** *Let t be a non-negative integer. For any t -locally easy H , MIS can be $OPT^{1 - \frac{1}{t+1}}$ -approximated on H -free graphs.*

As we will see in Section 4.3 for forbidden graphs H containing a triangle, it is unlikely to improve the approximation ratio below $n^{1/4}$. However if H is a star, constant-approximations are achievable. It is known that MIS is polynomial-time solvable on claw-free graphs [31] (i.e., $K_{1,3}$ -free graphs) while it is APX-hard on $K_{1,4}$ -free graphs (see for instance [3]). The greedy algorithm (or actually any sensible algorithm) gives an s -approximation in $K_{1,s}$ -free graphs. The ratio was improved to arbitrarily close to $\frac{s-1}{2}$ by Halldórsson.

► **Theorem 17** ([25]). *For every $s \geq 4$ and $\varepsilon > 0$, MIS is $\frac{s-1}{2} + \varepsilon$ -approximable on $K_{1,s}$ -free graphs.*

3.3 Local Search for $K_{t,t}$ -free graphs

Here we analyze the performance of the t -LOCAL SEARCH algorithm in $K_{t,t}$ -free graphs. Usually for the particular case of the biclique, “ $K_{t,t}$ -free” is intended as “no $K_{t,t}$ as a subgraph”. Here we still mean “no $K_{t,t}$ as an induced subgraph”, since our algorithm works even in this more general setting. For a fixed integer $t > 2$, t -LOCAL SEARCH takes as input a graph G , and constructs an independent set S from a single vertex. Then, it tries to improve S in the following way: whenever there exist two sets $X \subseteq S$ (note that X can possibly be empty) and $Y \subseteq V \setminus S$ such that $0 \leq |X| < |Y| \leq t$ and $(S \setminus X) \cup Y$ is an independent set, it replaces S by $(S \setminus X) \cup Y$ (if there are several choices, it chooses an arbitrary one). When S can no longer be improved, it outputs it. Each improvement takes $O(n^{2t})$ time, and the number of such improvements is at most n , since the size of S increases by at least one at each step. Hence, the algorithm takes polynomial time. In the following theorem, we prove that this simple algorithm provides an $O(OPT^{1-1/t})$ -approximation whenever the input graph is $K_{t,t}$ -free. In particular, 2-LOCAL SEARCH is an $O(\sqrt{OPT})$ -approximation in C_4 -free graphs. It came to our knowledge that the same result was obtained independently by Dvořák, Feldmann, Rai, and Rzażewski [15].

► **Theorem 18.** *For any fixed $t \geq 2$, t -LOCAL SEARCH is an $O(OPT^{1-1/t})$ -approximation in $K_{t,t}$ -free graphs.*

Proof. Let S be the solution returned by the algorithm, and O be a fixed optimal solution. The objective is to bound $|O'|$ in terms of $|S'|$, where $O' := O \setminus S$ and $S' := S \setminus O$. To this end, let us consider $\mathcal{B} := G[S' \cup O']$ the bipartite graph induced by $S' \cup O'$. Let $k := |S'|$. We partition O' into D^- and D^+ , where D^- are the vertices of O' whose degree within S' is at most $t - 1$, and thus D^+ are the vertices of O' whose degree within S' is at least t . We now bound the sizes of D^- and D^+ separately.

- Let us partition D^- into classes D_1^-, \dots, D_q^- with respect to the equivalence relation $u \sim v$ if and only if $N_{S'}(u) = N_{S'}(v)$. By definition of D^- we have $q \leq \sum_{i=1}^{t-1} \binom{k}{i}$. Then, we claim that for every $i \in \{1, \dots, q\}$, we have $|D_i^-| \leq t - 1$. Indeed, we must have $|D_i^-| \leq |N_{S'}(D_i^-)|$, since otherwise the algorithm would have replaced S by $(S \setminus N_{S'}(D_i^-)) \cup D_i^-$. This proves $|D^-| \leq (t - 1) \sum_{i=1}^{t-1} \binom{k}{i}$.

- For a set $X \subseteq S'$, let $I_X := \bigcap_{x \in X} N_{O'}(x)$. Observe that if $|X| = t$, then necessarily $I_X \subseteq D^+$, and moreover $|I_X| \leq t - 1$, since otherwise the graph would have an induced $K_{t,t}$. Finally, we have $D^+ = \bigcup_{X \subseteq S', |X|=t} I_X$, which proves that $|D^+| \leq \binom{k}{t}(t - 1)$.

Hence we have $|O'| \leq |D^-| + |D^+| \leq (t - 1) \sum_{i=1}^{t-1} \binom{k}{i} + \binom{k}{t}(t - 1) = O(k^t)$. ◀

4 Graphs without short cycles

In this section we show that the strong inapproximability of MIS in general graphs survives, albeit in a less severe form, on graphs without small cycles. More quantitatively, we show that for any positive integer γ , there is a constant $\beta = \Theta(1/\gamma)$ depending only on γ , such that an n^β -approximation of MIS in graphs with girth γ is unlikely.

4.1 Triangle-free graphs

While Lemma 13 implies an $n^{1/2}$ -approximation of MIS in triangle-free graphs, a natural question is how much the ratio's exponent can be decreased. In this section we provide a lower bound for it.

The following result will be made obsolete twice. Indeed we will then generalize its statement from triangle-free, that is girth 4, to graphs with any constant girth. Then in Section 4.3 we will present a stronger inapproximability result of $\Omega(n^{1/4-\varepsilon})$. Nevertheless we choose to keep its proof as it is simpler, easier to follow, and self-contained. Furthermore, it contains all the ideas necessary to achieve the subsequent results.

► **Theorem 19.** *For any $\varepsilon > 0$, it is NP-hard to distinguish between triangle-free graphs G on n vertices satisfying*

- $\alpha(G) \leq n^{5/6-\varepsilon}$, and
- $\alpha(G) \geq n^{1-\varepsilon}$.

So for any $\varepsilon > 0$, MIS cannot be approximated within ratio $n^{1/6-\varepsilon}$ in triangle-free graphs unless $NP \subseteq BPP$.

Proof. Let $\varepsilon > 0$ be an arbitrarily small real value, and $\varepsilon := 3\varepsilon$. We perform a randomized reduction from an infinite set of graphs H admitting the following gap: Positive instances have stable sets of size at least $|V(H)|^{1-\varepsilon}$ whereas negative instances have no stable set of size $|V(H)|^\varepsilon$. It is known that distinguishing between these two cases is NP-hard for randomized reductions [27], and even for deterministic ones [35].

Reduction. Given an N -vertex graph H , we construct a triangle-free graph G in the following way. We transform every vertex v of H into an independent set $I(v)$ of size $s := N^5$. For every edge $uv \in E(H)$, we put a random bipartite graph between $I(u)$ and $I(v)$: for each pair of vertices $x \in I(u)$, $y \in I(v)$, we independently add an edge xy to $E(G)$ with probability $p := N^{-4-\eta}$ with $\eta := 2\varepsilon/3$. We denote by G_Δ the graph thus obtained. A key property is that G_Δ contains only *few* triangles. For each triangle in G_Δ , we remove all three vertices of it. We call that phase the *triangle removal*, and we denote by G the triangle-free graph that arises when that phase comes to an end. We further assume that N is larger than the smallest integral constant N_0 for which for every $N \geq N_0$, $N^{3\eta} > N^{2.5\eta} + 10N^{2\eta} \ln N$, $N^{-\varepsilon} > 6N^{-2\eta}$, $2^{17/6}N^{\eta/3} < N^\varepsilon$, and $N^{-\eta} < 10^{-100}$. In particular the second and third inequalities hold for sufficiently large N since $\eta < \varepsilon = 3\varepsilon < 2\eta$. The hardness of approximation [27, 35] still holds since instances with less than a constant number of vertices can be solved optimally in constant time.

► **Lemma 20.** *For every edge $uv \in E(H)$, the probability that there exist two sets $A \subset I(u), B \subset I(v)$ both of size $N^{4+2\eta}$ without any edge between A and B is at most $e^{-N^{4+2.5\eta}}$. Thus, with high probability, this event does not happen.*

Proof. The probability that there is no edge between two fixed sets A and B of size $N^{4+2\eta}$ is:

$$(1-p)^{|A| \cdot |B|} = \left(1 - \frac{1}{N^{4+\eta}}\right)^{N^{2(4+2\eta)}} \leq e^{-N^{4+3\eta}}.$$

By the union bound, the probability that there is at least one such pair of sets is at most:

$$\left(\frac{N^5}{N^{4+2\eta}}\right)^2 e^{-N^{4+3\eta}} \leq N^{10N^{4+2\eta}} e^{-N^{4+3\eta}} = e^{-N^{4+3\eta} + 10N^{4+2\eta} \ln N} \leq e^{-N^{4+2.5\eta}}.$$

◀

► **Lemma 21.** *The expected number of triangles in G_Δ is at most $N^{6-3\eta}$. Furthermore $|V(G_\Delta)| - |V(G)|$ is at most $3N^{6-2\eta}$ with probability at least $1 - N^{-\eta}$.*

Proof. The expected number of triangles in G_Δ is:

$$\mathbb{E}(\#\langle \Delta, G_\Delta \rangle) \leq (sN)^3 p^3 = (N^6)^3 N^{-12-3\eta} = N^{6-3\eta}.$$

By Markov's inequality, $\mathbb{P}(\#\langle \Delta, G_\Delta \rangle \geq N^{6-2\eta}) \leq N^{6-3\eta} / N^{6-2\eta} = N^{-\eta}$.

◀

Let n be the number of vertices of G (after the triangle removal). By the previous lemma $n := |V(G)| > N^6/2$, with high probability.

If H is a YES-instance, there is a stable set of size $n^{1-\varepsilon}$ in G . We assume that H is a YES-instance, so there is a stable set S in H such that $|S| \geq N^{1-\varepsilon}$. By construction, $S_{G_\Delta} := \bigcup_{u \in S} I(u)$ is a stable set in G_Δ of size $s|S| \geq N^{6-\varepsilon}$. By Lemma 21, $S_{G_\Delta} \cap V(G)$ is an independent set in G of size, w.h.p., at least $N^{6-\varepsilon} - 3N^{6-2\eta} > N^{6-\varepsilon}/2 > n^{1-\varepsilon/6}/2 = n^{1-\varepsilon/2}/2 > n^{1-\varepsilon}$.

If H is a NO-instance, there is no stable set of size $n^{5/6+\varepsilon}$ in G . Let S_G be an independent set of G and let $S := \{v \in V(H) \text{ such that } |I(v) \cap S_G| \geq N^{4+2\eta}\}$. If H is a NO-instance, then there is no stable set in H of size more than N^ε . By a union bound of applications of Lemma 20 to all pairs of vertices of S , w.h.p S is an independent set of H , which implies that $|S| < N^\varepsilon$. Thus $|S_G| < sN^\varepsilon + N^{4+2\eta}(N - N^\varepsilon) < N^{5+\varepsilon} + N^{5+2\eta} < 2N^{5+2\eta} = 2N^{5(1+2\eta/5)} < 2^{11/6} n^{5/6+\eta/3} < n^{5/6+\varepsilon}$.

◀

4.2 Graphs with higher girth

Monien, Speckenmeyer and Murphy independently found improved approximations when the girth, actually even the odd girth, is any constant γ .

► **Theorem 22** ([32, 33]). *MIS admits a polynomial-time $n^{\frac{2}{\gamma-1}}$ -approximation on graphs with odd girth γ .*

In particular, the result implies an $n^{1/2}$ -approximation for triangle-free graphs, an $n^{1/3}$ -approximation for $\{C_3, C_5\}$ -free graphs, an $n^{1/4}$ -approximation for $\{C_3, C_5, C_7\}$ -free graphs, etc. On the complexity side, the construction of Theorem 19 where the probability p of having an edge between $I(u)$ and $I(v)$ with $uv \in E(H)$ is now set to $N^{-2(\gamma-1)-\eta}$ and the size s of each $I(u)$ is set to $N^{2\gamma-1}$ yields a polynomial gap on C_γ -free graphs, and even on graphs with girth $\gamma + 1$.

► **Theorem 23.** *For any $\varepsilon > 0$, it is NP-hard to distinguish between graphs G with n vertices and girth $\gamma + 1$ satisfying*

- $\alpha(G) \leq n^{\frac{2\gamma-1}{2\gamma}-\varepsilon}$, and
- $\alpha(G) \geq n^{1-\varepsilon}$.

Hence, for any $\varepsilon > 0$, MIS cannot be approximated within ratio $n^{\frac{1}{2\gamma}-\varepsilon}$ in graphs with girth $\gamma + 1$ unless $NP \subseteq BPP$.

Proof. We do the same reduction as in Theorem 19 with the following modifications. We now set $\varepsilon := \gamma\varepsilon$, $s := N^{2\gamma-1}$, $p := N^{2(\gamma-1)-\eta}$, and $\eta := \frac{\gamma-1}{\gamma}\varepsilon$. We denote by G_o the graph obtained before the removal step. For every cycle of length at most γ , we remove all the vertices of the cycle from the graph. When this short cycle removal ends, the graph has girth at least $\gamma + 1$. We call G the obtained graph.

► **Lemma 24.** *For every edge $uv \in E(H)$, the probability that there exist two sets $A \subset I(u), B \subset I(v)$ both of size $N^{2(\gamma-1)+2\eta}$ without any edge between A and B is at most $e^{-N^{2(\gamma-1)+2.5\eta}}$. Thus, with high probability, this event does not happen.*

Proof. The probability that there is no edge between two fixed sets A and B of size $N^{2(\gamma-1)+2\eta}$ is:

$$(1-p)^{|A| \cdot |B|} = \left(1 - \frac{1}{N^{2(\gamma-1)+\eta}}\right)^{N^{2(2(\gamma-1)+2\eta)}} \leq e^{-N^{2(\gamma-1)+3\eta}}.$$

By the union bound, the probability that there is at least one such pair of sets is at most:

$$\begin{aligned} \left(\frac{N^{2\gamma-1}}{N^{2(\gamma-1)+2\eta}}\right)^2 e^{-N^{2(\gamma-1)+3\eta}} &\leq N^{10N^{2(\gamma-1)+2\eta}} e^{-N^{2(\gamma-1)+3\eta}} \\ &= e^{-N^{2(\gamma-1)+3\eta} + 10N^{2(\gamma-1)+2\eta} \ln N} \leq e^{-N^{2(\gamma-1)+2.5\eta}}. \end{aligned}$$

◀

► **Lemma 25.** *The expected number of cycles of length at most γ in G_o is at most $N^{(2-\eta)\gamma}$. Furthermore $|V(G_o)| - |V(G)|$ is at most $\gamma N^{2\gamma-\eta(\gamma-1)}$ with probability at least $1 - N^{-\eta}$.*

Proof. The expected number of cycles of length at most γ in G_o is:

$$\mathbb{E}(\#(C_{3 \rightarrow \gamma}, G_o)) \leq \gamma (sN)^\gamma p^\gamma = \gamma N^{2\gamma^2} N^{-(2(\gamma-1)-\eta)\gamma} = \gamma N^{(2-\eta)\gamma}.$$

By Markov's inequality, $\mathbb{P}(\#(C_{3 \rightarrow \gamma}, G_o) \geq \gamma N^{2\gamma-\eta(\gamma-1)}) \leq N^{(2-\eta)\gamma} / \gamma N^{2\gamma-\eta(\gamma-1)} = N^{-\eta}$.

◀

Let n be the number of vertices of G (after the short cycle removal). By the previous lemma $n := |V(G)| > N^{2\gamma}/2$, with high probability.

If H is a YES-instance, there is a stable set of size $n^{1-\varepsilon}$ in G . We assume that H is a YES-instance, so there is a stable set S in H such that $|S| \geq N^{1-\varepsilon}$. By construction, $S_{G_o} := \bigcup_{u \in S} I(u)$ is a stable set in G_o of size $s|S| \geq N^{2\gamma-\varepsilon}$. By Lemma 25, $S_{G_o} \cap V(G)$ is an independent set in G of size, w.h.p., at least $N^{2\gamma-\varepsilon} - \gamma N^{2\gamma-\eta(\gamma-1)} > N^{2\gamma-\varepsilon}/2 > n^{1-\varepsilon/2\gamma}/2 = n^{1-\varepsilon/2}/2 > n^{1-\varepsilon}$.

If H is a NO-instance, there is no stable set of size $n^{(2\gamma-1)/(2\gamma)+\varepsilon}$ in G . Let S_G be an independent set of G and let $S := \{v \in V(H) \text{ such that } |I(v) \cap S_G| \geq N^{2(\gamma-1)+2\eta}\}$. If H is a NO-instance, then there is no stable set in H of size more than N^ε . By a union bound of

applications of Lemma 24 to all pairs of vertices of S , w.h.p S is an independent set of H , which implies that $|S| < N^\varepsilon$. Thus $|S_G| < sN^\varepsilon + N^{2(\gamma-1)+2\eta}(N - N^\varepsilon) < N^{2\gamma-1+\varepsilon} + N^{2\gamma-1+2\eta} < 2N^{2\gamma-1+2\eta} = 2N^{(2\gamma-1)(1+2\eta/(2\gamma-1))} < 2^{(4\gamma-1)/(2\gamma)} n^{(2\gamma-1)/(2\gamma)+\eta/\gamma} < n^{(2\gamma-1)/(2\gamma)+\varepsilon}$. ◀

Let us note that there is still a 4-fold multiplicative factor in the exponent between the approximation ratios of Theorem 22, namely $n^{2/\gamma}$, and the hardness ratios of $n^{1/(2\gamma)-o(1)}$ in Theorem 23. It is an interesting open question to bridge this gap.

An even hole is an induced cycle of even length at least 4. Even-hole-free graphs are $\{C_4, C_6, C_8, \dots\}$ -free graphs. The computational complexity of MIS on even-hole-free graphs is still unknown. An FPT algorithm was established recently [26]. We observe that Local Search readily gives a PTAS for that problem. We leave the existence of an EPTAS as an open problem.

► **Observation 26.** MIS can be $(1+\varepsilon)$ -approximated in time $n^{O(1/\varepsilon)}$ on even-hole-free graphs.

Proof. The graph induced by the symmetric difference between any two feasible solutions is bipartite and even-hole-free, hence it is a forest. Let S be a solution obtained by local search on an input graph G , O be a fixed optimum solution, $S' := S \setminus O$, and $O' := O \setminus S$. It is known that when $G[S' \cup O']$ is planar, there is an absolute constant C such that the C/ε^2 -LOCAL SEARCH $1 + \varepsilon$ -approximates the problem [34, 11], that is for a maximization problem, $|S'| \geq (1 - \varepsilon)|O'|$, implying $|S| \geq (1 - \varepsilon)|O|$. This gives a PTAS with running time $n^{O(1/\varepsilon^2)}$. When $G[S' \cup O']$ is even a forest, then it can be shown, and it is somewhat folklore, that a C/ε -LOCAL SEARCH is sufficient. ◀

4.3 Strengthening the inapproximability

There are two directions to improve the hardness-of-approximation results of Sections 4.1 and 4.2. As already mentioned, one can try to match upper and lower bounds in the approximation ratio, or at least to increase the exponent δ such that an n^δ -approximation would contradict a standard complexity-theoretic assumption. For triangle-free graphs, for instance, we do *not* expect a matching $n^{1/6}$ -approximation. And a likely outcome is that, ignoring logarithmic factors, the \sqrt{n} -approximation is best possible. We will actually show that an $n^{1/4-\varepsilon}$ -approximation is unlikely. The other direction is to derandomize our reductions. That way the inapproximability would be subject to the more (arguably the most) standard complexity assumption that P is not equal to NP. Derandomizing without degrading the quality of the gap seems challenging. We now encapsulate the reductions of Sections 4.1 and 4.2 so that both improving tasks boil down to exhibiting a randomized or deterministic family of graphs.

We say that an infinite family of graphs \mathcal{C} is *non-disappearing* if there is a constant $K \in (0, 1]$ such that for every positive integer n , there is a graph $G \in \mathcal{C}$ with at least Kn and at most n/K vertices. A non-disappearing family is called *efficient* if there is a polynomial-time algorithm which given an integer n (encoded in unary), outputs such a graph G . For example, a family containing at least one graph for every number of vertices is non-disappearing. We denote by \mathcal{G}_γ the set of all graphs with girth at least γ .

► **Theorem 27.** Let $\gamma > 3$ be an integer, $\delta \in (0, 1)$ be a real (allowed to depend on γ), and \mathcal{C} be an efficient non-disappearing family included in \mathcal{G}_γ , such that for every $G \in \mathcal{C}$ there is no disjoint pair of sets $A, B \subseteq V(G)$ satisfying both $|A| = |B| \geq |V(G)|^\delta$ and $E(A, B) = \emptyset$. Then MIS in \mathcal{G}_γ cannot be $n^{\frac{1-\delta}{2}-\varepsilon}$ -approximated, unless $P = NP$.

Proof. We assume all the preconditions hold and follow the construction of Theorems 19 and 23. We again draw a graph F from graphs of size N and gap $N^{1-\varepsilon}$. We substitute every vertex v by an independent set $I(v)$ of size between $KN^{\frac{1+\delta}{1-\delta}}$ and $\frac{1}{K}N^{\frac{1+\delta}{1-\delta}}$ such that there is a $G \in \mathcal{C}$ of the same size as the obtained graph G' , and the sets $I(v)$ are balanced (their size differs by at most 1). Both graphs have $\Theta(N^{\frac{1+\delta}{1-\delta}+1}) = \Theta(N^{\frac{2}{1-\delta}})$ vertices, say $cN^{\frac{2}{1-\delta}}$. We arbitrarily identify the vertices of G and G' in a one-to-one mapping. We keep an edge between two vertices u and v if uv is both an edge in G and G' . Thus we do the “intersection” of G and G' . We call J the final result.

Since G has girth at least γ , J has also girth at least γ . By assumption on \mathcal{C} , if there is an edge between u and v in F , then for every $A \subset I(u)$ and $B \subset I(v)$ both of size $c^\delta N^{\frac{2\delta}{1-\delta}}$, there is at least one edge in $E_J(A, B)$. We observe that $N^{\frac{2\delta}{1-\delta}} = N^{\frac{1+\delta}{1-\delta}-1}$ which is, up to constant multiplicative factors, the size of an $I(w)$ divided by N . Therefore we have the same important property as in Theorems 19 and 23. Thus we can finish the proof similarly, and conclude that distinguishing between instances with independence number at most $N^{\frac{1+\delta}{1-\delta}+\varepsilon}$ or at least $N^{\frac{2}{1-\delta}-\varepsilon}$ is NP-hard, for an arbitrary small $\varepsilon > 0$. The gap is $N^{1-\varepsilon}$ and $n := |V(J)| = \Theta(N^{\frac{2}{1-\delta}})$, hence a gap of $n^{\frac{1-\delta}{2}-\varepsilon}$. ◀

We now give a randomized counterpart of the previous theorem. For any integer $\gamma > 3$ and real $\delta \in (0, 1)$, we say that a distribution of graphs \mathcal{D} is (γ, δ) -appropriate if there is a constant $K \in (0, 1]$ and a polynomial-time algorithm, that given an integer n (encoded in unary), draws a graph G of size at least Kn and at most n/K out of this distribution such that with high probability, G has girth at least γ and no disjoint pair of sets $A, B \subseteq V(G)$ satisfies both $|A| = |B| \geq |V(G)|^\delta$ and $E(A, B) = \emptyset$.

► **Theorem 28.** *Let $\gamma > 3$ be an integer, $\delta \in (0, 1)$ be a real (allowed to depend on γ), and \mathcal{D} be a (γ, δ) -appropriate distribution. Then MIS in \mathcal{G}_γ cannot be $n^{\frac{1-\delta}{2}-\varepsilon}$ -approximated, unless $NP \subseteq BPP$.*

Proof. The proof is the same as Theorem 27, using a graph drawn from the distribution \mathcal{D} instead of a deterministic one from \mathcal{C} . Therefore we need the stronger assumption that NP is not contained in BPP. ◀

There are many constructions, all randomized, of triangle-free graphs with smallest possible independence number $\tilde{O}(\sqrt{n})$ [16, 29, 17, 28, 7]. These constructions all follow a simple scheme of starting from the empty graph, ordering the edges of the clique K_n , and then inserting an edge if it does not create a triangle, either among the inserted edges or among all the previous edges. The real difficulty is in the analysis of this probabilistic experiment. The logarithmic or constant factors were improved and the proofs simplified until Kim obtained a matching bound of $O(\sqrt{n \log n})$ [28]. This can be seen as the lower bound of $\Omega(n^2/\log n)$ for the off-diagonal Ramsey number $R(3, n)$, matching the upper bound $O(n^2/\log n)$ of Ajtai et al. [1].

To apply Theorem 28, we would need to check that the triangle-free graphs built in the aforementioned papers do not contain the complement of a large biclique K_{n^δ, n^δ} . As, for our purposes, we do not need the optimal bound of Kim, we follow the original proof of Erdős [16] giving the bound of $O(\sqrt{n \log n})$. Going through all the lemmas and replacing occurrences of K_x , where $x = O(\sqrt{n \log n})$, by $K_{x,x}$, the desired result can be obtained. In our language, the process described in the previous paragraph yields a $(4, 1/2)$ -appropriate distribution. This together with Theorem 28 improves the inapproximability of Theorem 19.

► **Corollary 29.** *MIS in \mathcal{G}_4 (i.e., triangle-free graphs) cannot be $n^{\frac{1}{4}-\varepsilon}$ -approximated, unless $NP \subseteq BPP$.*

We are not aware of any explicit deterministic construction of triangle-free graphs whose complements do not contain $K_{n^{2/3}, n^{2/3}}$ as a subgraph (which would derandomize Theorem 19), let alone, $K_{\sqrt{n}, \sqrt{n}}$. Deterministic constructions of graphs with large girth and large chromatic number, such as Ramanujan graphs with non-constant degree, might give some lower bound via Theorem 27, but not as good as Corollary 29. Actually, being based on a tight construction, the inapproximability of Corollary 29 can only be improved via a totally different route. One should also not completely rule out that there is an $n^{1/4}$ -approximation for MIS on triangle-free graphs.

5 Concluding remarks

The Erdős-Hajnal conjecture has proven particularly difficult. For example, the cases of P_5 -free or C_5 -free graphs are both wide open. For the few graphs H for which a proof that the Erdős-Hajnal property holds, it appears that the proof comes with an efficient algorithm reporting a sufficiently large independent set or clique. This is what we called the constructive Erdős-Hajnal property. We proposed a first and more humble step (see Theorem 6) in proving that a graph H has the constructive Erdős-Hajnal property: show that MIS in H -free graphs can be approximated within ratio $n^{1-\varepsilon}$ for an $\varepsilon > 0$, an unachievable ratio in general graphs. As mentioned in the introduction, this is strictly simpler than Erdős-Hajnal considering the case of P_5 . Yet it does not seem to us that this weaker conjecture is that much simpler now considering the graph C_5 . We believe that efforts to settle the improved approximation conjecture might turn out useful to make progress on the Erdős-Hajnal conjecture. In general, a cross-fertilization between Approximability Theory and the study of favorable Ramsey properties may prove fruitful. In particular, obtaining an $n^{0.99}$ -approximation algorithm for MIS in C_5 -free seems like a challenging open question.

Of course, classifying the approximability of MAXIMUM INDEPENDENT SET in H -free graphs is also an interesting task by its own means. On the one hand, already known reductions rule out PTASes in most H -free graphs classes, namely for any connected H different from a path or a subdivision of a claw. On the other hand, a constant-approximation algorithm can be turned into a PTAS in many H -free classes, by running the approximation on the input graph elevated to some appropriate power (using for instance the lexicographic product). This trick, originally used to rule out approximation algorithms for MAX CLIQUE in general graphs [20], works in the setting of H -free classes when H satisfies some properties, such as being a prime graph (i.e., having no non-trivial module). Hence, although MIS admits a constant-factor approximation in $K_{1,t}$ -free graphs for any $t \in \mathbb{N}$ (as mentioned in Section 3), it is not in APX when the forbidden graph is a simple tree, such as the 1-subdivision of $K_{1,4}$. Finally, another interesting consequence of the previous observation concerns P_t -free graphs: any constant-factor approximation for MIS in P_t -free graphs implies a PTAS (notice that the current “best” approximation algorithm in P_t -free graphs is a quasi-polynomial approximation scheme [13]).

Acknowledgment. We would like to thank Colin Geniet for pointing out to us the remark on graph products mentioned in the conclusion.

References

- 1 Miklós Ajtai, János Komlós, and Endre Szemerédi. A note on ramsey numbers. *J. Comb. Theory, Ser. A*, 29(3):354–360, 1980. URL: [https://doi.org/10.1016/0097-3165\(80\)90030-8](https://doi.org/10.1016/0097-3165(80)90030-8), doi:10.1016/0097-3165(80)90030-8.

- 2 S. Alekseev. A note on stable sets and colorings of graphs. *Commentationes Mathematicae Universitatis Carolinae*, 15, issue 2:307–309, 1974.
- 3 V. E. Alekseev. The effect of local constraints on the complexity of determination of the graph independence number. *Combinatorial-algebraic methods in applied mathematics*, pages 3–13, 1982.
- 4 Vladimir E. Alekseev. Polynomial algorithm for finding the largest independent sets in graphs without forks. *Discrete Applied Mathematics*, 135(1-3):3–16, 2004. URL: [https://doi.org/10.1016/S0166-218X\(02\)00290-1](https://doi.org/10.1016/S0166-218X(02)00290-1), doi:10.1016/S0166-218X(02)00290-1.
- 5 Paola Alimonti and Viggo Kann. Some apx-completeness results for cubic graphs. *Theoretical Computer Science*, 237(1):123 – 134, 2000. URL: <http://www.sciencedirect.com/science/article/pii/S0304397598001583>, doi:[https://doi.org/10.1016/S0304-3975\(98\)00158-3](https://doi.org/10.1016/S0304-3975(98)00158-3).
- 6 Brenda S. Baker. Approximation algorithms for np-complete problems on planar graphs. *J. ACM*, 41(1):153–180, January 1994. URL: <http://doi.acm.org/10.1145/174644.174650>, doi:10.1145/174644.174650.
- 7 Béla Bollobás. *Random Graphs, Second Edition*, volume 73 of *Cambridge Studies in Advanced Mathematics*. Cambridge University Press, 2011. URL: <https://doi.org/10.1017/CB09780511814068>, doi:10.1017/CB09780511814068.
- 8 John Adrian Bondy and Uppaluri S. R. Murty. *Graph Theory*. Graduate Texts in Mathematics. Springer, 2008. URL: <https://doi.org/10.1007/978-1-84628-970-5>, doi:10.1007/978-1-84628-970-5.
- 9 Andreas Brandstädt and Raffaele Mosca. Maximum weight independent set for claw-free graphs in polynomial time. *Discrete Applied Mathematics*, 237:57–64, 2018. URL: <https://doi.org/10.1016/j.dam.2017.11.029>, doi:10.1016/j.dam.2017.11.029.
- 10 Timothy M. Chan. Polynomial-time approximation schemes for packing and piercing fat objects. *J. Algorithms*, 46(2):178–189, 2003. URL: [https://doi.org/10.1016/S0196-6774\(02\)00294-8](https://doi.org/10.1016/S0196-6774(02)00294-8), doi:10.1016/S0196-6774(02)00294-8.
- 11 Timothy M. Chan and Sariel Har-Peled. Approximation algorithms for maximum independent set of pseudo-disks. *Discrete & Computational Geometry*, 48(2):373–392, 2012. URL: <https://doi.org/10.1007/s00454-012-9417-5>, doi:10.1007/s00454-012-9417-5.
- 12 Maria Chudnovsky. The erdős-hajnal conjecture - A survey. *Journal of Graph Theory*, 75(2):178–190, 2014. URL: <https://doi.org/10.1002/jgt.21730>, doi:10.1002/jgt.21730.
- 13 Maria Chudnovsky, Marcin Pilipczuk, Michal Pilipczuk, and Stéphan Thomassé. Quasi-polynomial time approximation schemes for the maximum weight independent set problem in H -free graphs. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*, pages 2260–2278, 2020. URL: <https://doi.org/10.1137/1.9781611975994.139>, doi:10.1137/1.9781611975994.139.
- 14 E. D. Demaine, M. T. Hajiaghayi, and K. Kawarabayashi. Algorithmic graph minor theory: Decomposition, approximation, and coloring. In *46th Annual IEEE Symposium on Foundations of Computer Science (FOCS'05)*, pages 637–646, Oct 2005. doi:10.1109/SFCS.2005.14.
- 15 Pavel Dvořák, Andreas Emil Feldmann, Ashutosh Rai, and Paweł Rzażewski. private communication. 2019.
- 16 Paul Erdős. Graph theory and probability. ii. *Canadian Journal of Mathematics*, 13:346–352, 1961.
- 17 Paul Erdős, Stephen Suen, and Peter Winkler. On the size of a random maximal graph. *Random Struct. Algorithms*, 6(2/3):309–318, 1995. URL: <https://doi.org/10.1002/rsa.3240060217>, doi:10.1002/rsa.3240060217.
- 18 P. Erdős and A. Hajnal. Ramsey-type theorems. *Discrete Applied Mathematics*, 25(1):37 – 52, 1989. URL: <http://www.sciencedirect.com/science/article/pii/0166218X89900450>, doi:[https://doi.org/10.1016/0166-218X\(89\)90045-0](https://doi.org/10.1016/0166-218X(89)90045-0).
- 19 P. Erdős. Some remarks on the theory of graphs. *Bulletin of the American Mathematical Society*, 53(4):292–294, 1947. doi:10.1090/S0002-9904-1947-08785-1.

- 20 Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. Approximating Clique is Almost NP-Complete (Preliminary Version). In *32nd Annual Symposium on Foundations of Computer Science, San Juan, Puerto Rico, 1-4 October 1991*, pages 2–12. IEEE Computer Society, 1991. URL: <https://doi.org/10.1109/SFCS.1991.185341>, doi:10.1109/SFCS.1991.185341.
- 21 Michael R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- 22 Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer, 1988. URL: <https://doi.org/10.1007/978-3-642-97881-4>, doi:10.1007/978-3-642-97881-4.
- 23 Andrzej Grzesik, Tereza Klimosova, Marcin Pilipczuk, and Michal Pilipczuk. Polynomial-time algorithm for maximum weight independent set on $\$p_6$ -free graphs. *CoRR*, abs/1707.05491, 2017. URL: <http://arxiv.org/abs/1707.05491>, arXiv:1707.05491.
- 24 Andrzej Grzesik, Tereza Klimosova, Marcin Pilipczuk, and Michal Pilipczuk. Polynomial-time algorithm for maximum weight independent set on P_6 -free graphs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, San Diego, California, USA, January 6-9, 2019*, pages 1257–1271, 2019. URL: <https://doi.org/10.1137/1.9781611975482.77>, doi:10.1137/1.9781611975482.77.
- 25 Magnús M. Halldórsson. Approximating discrete collections via local improvements. In *Proceedings of the Sixth Annual ACM-SIAM Symposium on Discrete Algorithms, 22-24 January 1995. San Francisco, California, USA.*, pages 160–169, 1995. URL: <http://dl.acm.org/citation.cfm?id=313651.313687>.
- 26 Edin Husić, Stéphan Thomassé, and Nicolas Trotignon. The independent set problem is FPT for even-hole-free graphs. *to appear in the proceedings of IPEC 2019*, 2019.
- 27 Johan Håstad. Clique is hard to approximate within $n^{1-\epsilon}$. In *Acta Mathematica*, pages 627–636, 1996.
- 28 Jeong Han Kim. The ramsey number $r(3, t)$ has order of magnitude $t^2/\log t$. *Random Structures & Algorithms*, 7(3):173–207, 1995.
- 29 Michael Krivelevich. Bounding ramsey numbers through large deviation inequalities. *Random Struct. Algorithms*, 7(2):145–156, 1995. URL: <https://doi.org/10.1002/rsa.3240070204>, doi:10.1002/rsa.3240070204.
- 30 Vadim V. Lozin and Martin Milanic. A polynomial algorithm to find an independent set of maximum weight in a fork-free graph. *J. Discrete Algorithms*, 6(4):595–604, 2008. URL: <https://doi.org/10.1016/j.jda.2008.04.001>, doi:10.1016/j.jda.2008.04.001.
- 31 George J. Minty. On maximal independent sets of vertices in claw-free graphs. *J. Comb. Theory, Ser. B*, 28(3):284–304, 1980. URL: [https://doi.org/10.1016/0095-8956\(80\)90074-X](https://doi.org/10.1016/0095-8956(80)90074-X), doi:10.1016/0095-8956(80)90074-X.
- 32 Burkhard Monien and Ewald Speckenmeyer. Ramsey numbers and an approximation algorithm for the vertex cover problem. *Acta Inf.*, 22(1):115–123, 1985. URL: <https://doi.org/10.1007/BF00290149>, doi:10.1007/BF00290149.
- 33 Owen J Murphy. Computing independent sets in graphs with large girth. *Discrete Applied Mathematics*, 35(2):167–170, 1992.
- 34 Nabil H. Mustafa and Saurabh Ray. Improved results on geometric hitting set problems. *Discrete & Computational Geometry*, 44(4):883–895, 2010. URL: <https://doi.org/10.1007/s00454-010-9285-9>, doi:10.1007/s00454-010-9285-9.
- 35 David Zuckerman. Linear degree extractors and the inapproximability of max clique and chromatic number. *Theory of Computing*, 3(1):103–128, 2007. URL: <https://doi.org/10.4086/toc.2007.v003a006>, doi:10.4086/toc.2007.v003a006.