Average Case Polyhedral Complexity of the Maximum Stable Set Problem

Gábor Braun[1], Samuel Fiorini[2], and Sebastian Pokutta[1]

1   ISyE, Georgia Institute of Technology, Atlanta, GA, USA
    `{gabor.braun,sebastian.pokutta}@isye.gatech.edu`
2   Department of Mathematics, Université libre de Bruxelles CP 216, Bd. du
    Triomphe, 1050 Brussels, Belgium
    `sfiorini@ulb.ac.be`

Abstract

We study the minimum number of constraints needed to formulate random instances of the maximum stable set problem via LPs (more precisely, linear extended formulations), in two distinct models. In the uniform model, the constraints of the LP are not allowed to depend on the input graph, which should be encoded solely in the objective function. There we prove a $2^{\Omega(n/\log n)}$ lower bound with probability at least $1-2^{-2^n}$ for every LP that is exact for a randomly selected set of instances; each graph on at most n vertices being selected independently with probability $p \geqslant 2^{-\binom{n/4}{2}+n}$. In the non-uniform model, the constraints of the LP may depend on the input graph, but we allow weights on the vertices. The input graph is sampled according to the $G(n,p)$ model. There we obtain upper and lower bounds holding with high probability for various ranges of p. We obtain a super-polynomial lower bound all the way from $p = \Omega(\frac{\log^{6+\varepsilon} n}{n})$ to $p = O(\frac{1}{\log n})$. Our upper bound is close to this as there is only an essentially quadratic gap in the exponent, which also exists in the worst-case model. Finally, we state a conjecture that would close this gap, both in the average-case and worst-case models.

1   Introduction

In the last three years, extended formulations considerably gained interest in various areas, including discrete mathematics, combinatorial optimization, and theoretical computer science. The key idea underlying extended formulations is that with the right choice of variables, various combinatorial optimization problems can be *efficiently* expressed via linear programs (LPs). This asks for the intrinsic difficulty of expressing optimization problems through a single LP, in terms of the minimum number of necessary *constraints*. This leads to a complexity measure that we call loosely here 'polyhedral complexity' (precise definitions are given later in Section 2).

On the one hand, there is an ever expanding collection of examples of small size extended formulations. For instance, [16] has expressed the minimum spanning tree problem on a planar graph with only a linear number of (variables and) constraints, while in the natural edge variables the LP has an exponential number of constraints. There exist numerous other examples, see e. g., the surveys by [7] and [11].

On the other hand, a recent series of breakthroughs in lower bounds renewed interest for extended formulations [14, 9, 2, 5, 3, 6, 15]. These breakthroughs make it now conceivable to quantify the polyhedral complexity of any given combinatorial optimization problem *unconditionally*, that is, independently of conjectures such as P vs. NP, and without extra assumption on the structure of the LP.

Although a polynomial upper bound on the polyhedral complexity yields a polynomial upper bound on the true algorithmic complexity of the problem—provided that the LP can be efficiently constructed and also that the size of the coefficients is kept under control (see [14] for a discussion of this last issue) e. g., through interior point methods—it is becoming clear that the converse does not hold. Recently, [6] proved that every LP for MAXCUT with approximation factor at most $2 - \varepsilon$ needs at least $n^{\Omega\left(\frac{\log n}{\log \log n}\right)}$ constraints, while the approximation factor of the celebrated SDP-based polynomial time algorithm of [10] is close to 1.13. Even more recently, [15] solved another major open problem in the area by showing a $2^{\Omega(n)}$ lower bound on the size of any LP expressing the perfect matching problem and in [4] it was shown that the matching polytope does not admit any fully-polynomial size relaxation scheme (the polyhedral equivalent of an FPTAS).

In this paper, we consider the problem of determining the *average case* polyhedral complexity of the maximum stable set problem, in two different models: 'uniform' and 'non-uniform', see Section 1.2 below. Roughly, the uniform model asks for a single LP that works for a given set of input graphs. In the non-uniform model the LP can depend on the input graph G but should work for every choice of weights on the vertices of G (in particular, for all induced subgraphs of G).

We show that the polyhedral complexity of the maximum stable set problem remains high in each of these models, when the input graph is sampled according to natural distributions. Therefore, we conclude that *the hardness of the maximum stable set problem is not concentrated on a small mass of graphs but is spread out through all graphs.*

1.1   Related Work

Our work is most directly related to [9] and [2], where the framework for bounding the size of approximate linear programming formulations was laid out. This framework forms the basis of our uniform model. We will also employ a robustness theorem from [3] for dropping constraints and feasible solutions.

1.2   Contribution

We present the first strong and unconditional results on the average case size of LP formulations for the maximum stable set problem. In particular, we establish that the maximum stable set problem in two natural average case models and encodings does not admit a polynomial size linear programming formulation, even in the unlikely case that P = NP.

Uniform Model

In the *uniform model* the polytope P containing the feasible solutions to the stable set problem is *independent* of the instances. The instances will be solely encoded into the objective functions. This ensures that no complexity of the problem is leaked into an instance-specific formulation. A good example of a uniform model is the TSP polytope over K_n with which we can test for Hamiltonian cycles in any graph with at most n vertices by choosing an appropriate objective function. In the uniform model, we consider the random

Figure 1 Comparing lower bounds on $\mathrm{xc}(G(n,p))$ for various regimes. For p close to $1/\sqrt[3]{n}$ the blue and green lines provide roughly the same bounds. For p significantly above $1/\sqrt[4]{n}$ the magenta line outperforms the green line.

instance set of graphs on at most n vertices, where each graph is contained in the instance set with probability $p \geqslant 2^{-\binom{n/4}{2}+n}$. Then we show that with probability at least $1 - 2^{-2^n}$, every LP formulation of this instance set in the natural encoding has at least $2^{\Omega(n/\log n)}$ constraints.

Non-uniform Model

In the *non-uniform model* we consider the stable set polytope for a *specific but random graph*. The polyhedral description may depend heavily on the chosen graph. We sample a graph G in the Erdős–Rényi $G(n,p)$ model, i. e., G has n vertices, and every pair of vertices is independently connected by an edge with probability p. We then analyze the stable set polytope $\mathrm{STAB}(G)$ of G. If p is small enough, so that the obtained graph is sufficiently sparse, it will contain an induced subgraph of sufficient size inducing a polyhedral reduction from the correlation polytope. Via this reduction we derive strong lower bounds on the size of any LP expressing $\mathrm{STAB}(G)$ that hold with high probability. In particular, we obtain superpolynomial lower bounds for p ranging between $\Omega(\frac{\log^{6+\varepsilon} n}{n})$ and $O(\frac{1}{\log n})$. For example for $p = n^{-\varepsilon}$ and $\varepsilon < 1/4$, any LP has at least $2^{\Omega(\sqrt{n^\varepsilon \log n})}$ constraints w.h.p., and for $p = \Omega(\frac{\log^{6+\varepsilon} n}{n})$, any LP has at least $n^{\log(3/2)\log^{\varepsilon/5} n}$ constraints w.h.p. Figure 1 illustrates our lower bounds.

1.3 Outline

In Section 2 we recall basics on extended formulations. We introduce the uniform model for the maximum stable set problem in Section 3.1. We then establish bounds on the average case complexity for the uniform model in Section 3. In Section 4 we consider the non-uniform model and derive lower bounds as well as upper bounds. We conclude with a conjecture in Section 5.

2 Preliminaries

We start by briefly recalling basics of extended formulations, stated in geometric terms. We refer the interested reader to [9] for more details. After that we state the main source of lower bounds in the non-uniform case.

Let $P \subseteq \mathbb{R}^d$ and $L \subseteq \mathbb{R}^e$ be two polyhedra. Then L is called an *extension* (or *lift*) of P if there exists an affine map $\pi \colon \mathbb{R}^d \to \mathbb{R}^e$, so that $\pi(L) = P$. Defining the *size* of polyhedron L as its number of facets, the *extension complexity* of polyhedron P is the minimum size of any of its extensions L, and is denoted by $\mathrm{xc}(P)$. Here we use the notions of extension and extended formulation interchangeably; the latter is simply an equivalent way to describe an extension.

The following monotonicity lemma from [9] provides a reduction mechanism to lower bound the extension complexity.

▶ **Lemma 1** (Monotonicity of extended formulations). *Let P be a polyhedron. Then the following hold:*
(i) *if F is a face of P, then $\mathrm{xc}(F) \leqslant \mathrm{xc}(P)$;*
(ii) *if L is an extension of P, then $\mathrm{xc}(P) \leqslant \mathrm{xc}(L)$.*

As usual, $\mathrm{corollary}(n) \coloneqq \mathrm{conv}\left(\{bb^\intercal \in \mathbb{R}^{n \times n} \mid b \in \{0,1\}^n\}\right)$ denotes the *correlation polytope* and $\mathrm{STAB}(G) \coloneqq \mathrm{conv}\left(\{\chi^S \in \mathbb{R}^{V(G)} \mid S \text{ stable set of } G\}\right)$ is the *stable set polytope* of graph G. (Recall that the characteristic vector χ^S has $\chi^S_v = 1$ if $v \in S$ and $\chi^S_v = 0$ otherwise.) Let log denote the base-2 logarithm.

▶ **Theorem 2.** $\mathrm{xc}(\mathrm{corollary}(n)) \geqslant 2^{n \cdot \log(3/2)}$.

The factor $\log(3/2) \approx 0.585$ in the exponent is the current best one due to [12]; for various approximate case versions see [2, 5, 3]. The first exponential lower bound was established in [9].

The notion of extension directly generalizes to pairs of nested polyhedra. If $P \subseteq Q \subseteq \mathbb{R}^d$ are two polyhedra, an extension of the pair P, Q is a polyhedron $L \subseteq \mathbb{R}^e$ such that $P \subseteq \pi(L) \subseteq Q$ for some affine map $\pi \colon \mathbb{R}^d \to \mathbb{R}^e$. The extension complexity $\mathrm{xc}(P, Q)$ of pair P, Q is the minimum size of an extension of that pair.

3 Average Case Complexity in the Uniform Model

3.1 A Uniform Model for Maximum Stable Set

Faithful linear encodings were introduced in [2] to study the polyhedral hardness of approximation of various problems, via pairs of polyhedra. Here we recall only the polyhedral pair arising from the standard encoding of the maximum stable set problem.

We take $P = \mathrm{corollary}(n)$ as inner polyhedron, which is informally the convex hull of all possible vertex subsets, just like the 0/1-cube, but this encoding allows additionally to count edges by a linear function. In fact, a potential stable set I is encoded as the correlation matrix $bb^\intercal$ plays, where b is the characteristic vector of I. The outer polyhedron $Q = Q(\mathcal{G})$ depends on a collection $\mathcal{G}$ of graphs with vertex set included in $[n]$, that defines the instances we wish to solve. For each graph G with $V(G) \subseteq [n]$, we define an objective function $w^G \in \mathbb{R}^{n \times n}$ encoding the maximum stable set problem: i. e., assigning its size to every stable set of G, and carefully extended to other vertex subsets, so that the maximum is still the maximal size of stable sets of G. To this aim, we disregard vertices outside G and penalize edges inside the vertex set. This is formally achieved by letting $w^G_{ij} = 1$ if $i = j$ and $i \in V(G)$, but $w^G_{ij} = w^G_{ji} = -1$ if $ij \in E(G)$, and finally $w^G_{ij} = 0$ otherwise. We define a polyhedron

$$Q(\mathcal{G}) \coloneqq \left\{ x \in \mathbb{R}^{n \times n}_+ \,\middle|\, \forall G \in \mathcal{G} \colon \langle w^G, x \rangle \leqslant \alpha(G) \right\},$$

of the constraints of maximum stable sets, where $\langle w^G, x \rangle = \sum_{i,j} w^G_{ij} x_{ij}$ denotes the Frobenius inner product of matrices w^G and x, and $\alpha(G)$ is the stability number of G. We let

$\text{STAB}^u(\mathcal{G}, \rho)$ denote the pair of nested polyhedra $(P, (1 + \rho)Q(\mathcal{G}))$ with $\rho \geqslant 0$ defining the dilation factor. If $\rho = 0$, we simply denote the pair by $\text{STAB}^u(\mathcal{G})$.

For each polyhedron K containing P and contained in $(1+\rho)Q(\mathcal{G})$, we have $\max\{\langle w^G, x \rangle \mid x \in K\} \geqslant \alpha(G)$ for all graphs G, so that K is a relaxation of the maximum stable set problem. Moreover, $\max\{\langle w^G, x \rangle \mid x \in K\} \leqslant (1+\rho)\alpha(G)$ for all $G \in \mathcal{G}$. In other words, the relaxation K is not more than a $(1 + \rho)$ factor off for all graphs in the collection $\mathcal{G}$.

The extension complexity of a polyhedral pair is equal to the nonnegative rank of any of its slack matrices up to a difference of 1, this is called the *factorization theorem*. For the pair $(P, (1 + \rho)Q(\mathcal{G}))$ a slack matrix S has rows indexed by all the characteristic vectors $b \in \{0,1\}^n$ of the subsets of $[n]$, corresponding to the vertex $bb^\intercal$ of P, and columns indexed by $\mathcal{G}$. The entries are $S(G, b) = (1 + \rho)\alpha(G) - \langle w^G, bb^\intercal \rangle$.

For example, when $\mathcal{G}$ is the set of all cliques, we may reindex the graphs by the characteristic vectors $a \in \{0,1\}^n$ of their vertex sets. We obtain a matrix M' as a slack matrix with rows and columns indexed by $a, b \in \{0,1\}^n$, and with entries $M'(a, b) = (1 - a^\intercal b)^2 + \rho$. Thus, in particular, restricting to the entries with $a^\intercal b \leq 1$, we obtain a partial matrix M

$$M(a, b) = \begin{cases} \rho & \text{if } a^\intercal b = 1 \\ 1 + \rho & \text{if } a^\intercal b = 0. \end{cases}$$

For $\rho = 0$, the partial matrix M is known as the *unique disjointness (UDISJ)* (partial) matrix. For general $\rho \geqslant 0$, this is called the *ρ-shifted* UDISJ matrix. We shall need the following theorem from [3] to bound the nonnegative rank of certain submatrices of the (ρ-shifted) UDISJ matrix.

▶ **Theorem 3.** *For the ρ-shifted UDISJ matrix M, let M_k be the submatrix for sets of size k. Let S be any submatrix of M_k obtained by deleting at most an α-fraction of rows and at most a β-fraction of columns for some $0 \leqslant \alpha, \beta < 1$. Then for $0 < \varepsilon < 1$:*

$$\text{rank}_+ S \geqslant 2^{(1/8(\rho+1) - (\alpha+\beta)\mathbb{H}[1/4])n - O(n^{1-\varepsilon})} \quad \text{for} \quad k = n/4 + O(n^{1-\varepsilon}).$$

3.2 Average Case Complexity

We will now establish our main result for the uniform average case complexity model. We obtain that for any random collection of graphs where each graph is picked independently with probability p, the polyhedral complexity of solving the stable set problem over that particular collection of graphs is high, or more precisely, the extension complexity of the corresponding pair is high. This shows in particular that the instances of the stable set problem resulting in high extension complexity are not localized in a set of small density.

▶ **Main Theorem 4** (Super-polynomial xc of $\text{STAB}^u(\mathcal{G})$ w.h.p.)**.** *Let $n \geqslant 40$ and $p \in [0, 1]$ with $p \geqslant 2^{-\binom{n/4}{2}+n}$. Pick a random family $\mathcal{G}$ of graphs by adding each graph G with $V(G) \subseteq [n]$ to the family with probability p, independent of the other G. Then*

$$\mathbb{P}\left[\text{xc}(\text{STAB}^u(\mathcal{G})) \geqslant 2^{\Omega(n/\log n)} \right] \geqslant 1 - 2^{-2^n}.$$

A crucial point of the proof is a concentration result on $\alpha(G)$. It is well-known that almost all graphs G on n vertices have stability number $\alpha(G) \sim 2 \log n$. However, the following rough estimate will be sufficient for our purpose, see e.g. [8, Proposition 11.3.4, page 304] for a proof.

▶ **Lemma 5.** *Let $n \geqslant 10$. The probability that a uniformly sampled random graph G with $V(G) = [n]$ has $\alpha(G) \geqslant 3 \log n$ is at most n^{-1}.*

We are ready to prove the main theorem of this section.

Proof of Main Theorem 4. The main idea of the proof is that, with large enough probability, we have $\max \left\{ \langle w^K, x \rangle \,\middle|\, x \in Q(\mathcal{G}) \right\} = O(\log n)$ for many cliques K with $V(K) \subseteq [n]$ and $\Theta(n)$ vertices. This implies that some slack matrix of the pair $\mathrm{STAB}^u(\mathcal{G})$ contains the $O(\log n)$-shifted UDISJ as a submatrix obtained by picking a large fraction of the rows (and all columns). We apply Theorem 3.

Consider a clique K with $V(K) \subseteq [n]$, and size $k := \lceil n/4 \rceil$. We say that a graph G is *good* for K if $V(G) = V(K)$ and $\alpha(G) \leqslant 3 \log n$. Clique K is said to be *good* if some graph $G \in \mathcal{G}$ is good for K. Otherwise, K is called *bad*.

We claim that, with high probability, the total fraction of bad cliques among all k-cliques K is at most $\alpha := 1/(24 \log n)$. By Lemma 5, the total number of graphs G with $V(G) = V(K)$ that are not good for a fixed k-clique K is at most $k^{-1} 2^{\binom{k}{2}}$. Thus

$$\mathbb{P}[K \text{ is bad}] = \mathbb{P}[\mathcal{G} \text{ contains no good graph for } K]$$

$$\leqslant (1-p)^{(1-k^{-1})2^{\binom{k}{2}}} \leqslant \mathrm{e}^{-p(1-k^{-1})2^{\binom{k}{2}}} \leqslant 2^{-\frac{9}{10} 2^n \log \mathrm{e}} \leqslant \alpha \, 2^{-2^n}.$$

where the second inequality follows from $k \geqslant n/4 \geqslant 10$ and $p \geqslant 2^{-\binom{n/4}{2}+n}$. Let X denote the random variable with value the number of bad k-cliques K. By Markov's inequality,

$$\mathbb{P}\left[X \geqslant \alpha \binom{n}{k} \right] \leqslant 2^{-2^n}.$$

If clique K is good and G is a good graph for K, the inequality $\langle w^G, x \rangle \leqslant 3 \log n$ is valid for $Q(\mathcal{G})$. Thus the inequality $\langle w^K, x \rangle \leqslant 3 \log n$ is also valid for $Q(\mathcal{G})$, because $x \geqslant 0$ is valid for $Q(\mathcal{G})$, and $w^K \leqslant w^G$.

Suppose that the fraction of cliques K with $V(K) \subseteq [n]$ and size $k = \lceil n/4 \rceil$ that are bad is at most α. We have shown that this holds with probability at least $1 - 2^{-2^n}$. By what precedes, we can define a slack matrix for the pair $\mathrm{STAB}^u(\mathcal{G})$ that contains a $(3 \log n)$-shift of UDISJ with at most an α-fraction of the rows thrown away. From Theorem 3 (with $\beta = 0$) and from the factorization theorem, the extension complexity of the pair $\mathrm{STAB}^u(\mathcal{G})$ is at least $2^{(1/8(3 \log n+1) - \alpha \mathbb{H}[1/4]) \cdot n - O(n^{1-\varepsilon})} = 2^{\Omega(n/\log n)}$. ◄

A close inspection of the proof of Main Theorem 4 shows that we can immediately apply the framework in [2] to obtain a lower bound on the average case *approximate* extension complexity. We obtain the following result. The proof is identical, except that we choose $\alpha := 1/24(1+\rho)\log n$, and the inequalities that yield the slack matrix are of the form $\langle w^K, x \rangle \leqslant (1+\rho)3 \log n$ for all good cliques K with $k = \lceil n/4 \rceil$ vertices, which are all valid for $(1+\rho)Q(\mathcal{G})$. A complete proof will be given in the full version of the paper.

▶ **Corollary 6** (Super-polynomial xc of $\mathrm{STAB}^u(\mathcal{G}, \rho)$ w.h.p.). *As in Main Theorem 4, let $\mathcal{G}$ be a random family of graphs such that each graph G with $V(G) \subseteq [n]$ is contained in $\mathcal{G}$ with probability $p \geqslant 2^{-\binom{n/4}{2}+n}$ independent of the other graphs. Then the ρ-approximate pair $\mathrm{STAB}^u(\mathcal{G}, \rho)$ with $\rho \leqslant \frac{n^{1-\varepsilon}}{\log n}$ for some $0 < \varepsilon < 1/2$ has extension complexity $2^{\Omega(n^\varepsilon)}$, with probability at least $1 - 2^{-2^n}$.*

Figure 2 Edge gadget E_{ij} replacing edge ij of T in the gadget graph $T^{\heartsuit}$. Black vertices represent vertices of the template graph T, white and grey vertices represent new vertices added to construct $T^{\heartsuit}$. There are $\ell \geqslant 0$ many edges between the black and grey vertices.

Observe that the approximation factor in Corollary 6 can be larger than $3 \log n$. The reason why this is possible, contradicting initial intuition, is that the hardness arises from having many different graphs and hence many objective functions to consider simultaneously and the encoding is highly non-monotone. Roughly speaking, graphs with different vertex sets are independent of each other, even if one is an induced subgraph of the other.

4 Average Case Complexity in the Non-uniform Model

We now turn our attention to the non-uniform model, where we consider the stable set polytope over a *specific but random* graph G and analyze its extension complexity. Our strategy is to embed certain gadget graphs as induced subgraphs of G, using the probabilistic method. Here we consider the Erdős–Rényi graph model and sample G from $G(n,p)$.

We begin by defining the gadget graphs we use and seeing how an induced gadget forces up the extension complexity of STAB(G).

Fix a graph T. This graph serves as a template for defining the *gadget graph* of T, denoted as $T^{\heartsuit}$: the graph obtained by replacing each edge ij of T with an *edge gadget* E_{ij}, which is a 5-cycle with two connecting paths (hairs) of length ℓ each, see Figure 2. In total, $T^{\heartsuit}$ has $v := |V(T)| + (2\ell + 3)|E(T)|$ vertices and $e := (2\ell + 5)|E(T)|$ edges. We will always take ℓ even but allow $\ell = 0$, in which case i and j are part of the 5-cycle (on Figure 2, both i and j then coincide with the closest gray vertex).

The hairs are used to decrease the average degree of induced subgraphs of the gadget graph, which makes it easier to embed $T^{\heartsuit}$ in $G(n,p)$ for lower values of p. This is formalized in the next lemma.

▶ **Lemma 7.** *For any graph T, the average degree of any induced subgraph of $T^{\heartsuit}$ is at most* $2 + 4/(2\ell + 3)$ *for* $\ell \geqslant 1$. *For* $\ell = 0$ *the average degree is at most 4.*

Proof. Let G be an induced subgraph of $T^{\heartsuit}$. We shall prove the stronger claim that $|E(G)| / |V(G) \setminus V(T)|$ is upper bounded by 2 if $\ell = 0$ and by $1 + 2/(2\ell + 3)$ if $\ell \geqslant 1$ (in other words, we ignore the original vertices of T at the estimation).

First we apply some modifications to G which do not decrease the factor $|E(G)| / |V(G) \setminus V(T)|$ if it was already at least 1. We add the original vertices of T to G (together with the edges connecting them to vertices already in G), and then we successively remove degree-1 vertices of G in the edge gadgets. Hence we may assume without loss of generality, that G has no degree-1 vertices of the edge gadgets, and it contains the original vertices of T. So for a fixed E_{ij}, the graph G contains either only the original vertices of T, or both the degree-3 (grey) vertices of the 5-cycle. In the latter case, from every path connecting these and i, j, the graph G contains either the whole path, or only the end points. We claim that if $\ell \geqslant 1$ then adding the missing paths will not decrease the factor $|E(G)| / |V(G) \setminus V(T)|$

below $1 + 2/(2\ell + 3)$ if it was greater than this value. Indeed, for every path, the ratio of added edges and vertices is at least $1 + 2/(2\ell + 3)$, namely, $1 + 1/(\ell - 1)$, 2 or 3/2. Therefore we may assume that every edge gadget E_{ij} is either completely contained in G or only the two vertices i and j of T are contained in G. Let k denote the number of E_{ij} completely contained in G, then $|E(G)| = k(2\ell + 5)$ and $|V(G) \setminus V(T)| = k(2\ell + 3)$, and their ratio is exactly $1 + 2/(2\ell + 3)$, finishing the proof in case $\ell \geqslant 1$.

If $\ell = 0$ then a similar argument applies, except that adding the shorter path (containing $u_{ij}^{(3)}$) and removing the longer path $(u_{ij}^{(1)}, u_{ij}^{(2)})$ will not decrease the factor $|E(G)| / |V(G) \setminus V(T)|$ below 2 if it were already larger. ◄

In the next lemma, we denote by corollary(T) the projection of the $|V(T)| \times |V(T)|$ correlation polytope corollary$(|V(T)|)$ on the variables x_{ii} for $i \in V(T)$ and x_{ij} for $ij \in E(T)$. We call this polytope the *correlation polytope of graph T*. In particular, corollary$(K_t) =$ corollary(t).

▶ **Lemma 8.** *If graph G contains $T^{\circlearrowleft}$ (with arbitrary even hair length ℓ) as an induced subgraph, then*

$$\mathrm{xc}(\mathrm{STAB}(G)) \geqslant \mathrm{xc}(\mathrm{corollary}(T)).$$

In particular, for $T = K_t$ we get $\mathrm{xc}(\mathrm{STAB}(G)) \geqslant 2^{t \log(3/2)}$.

Proof. Let F be the face of $\mathrm{STAB}(G)$ whose vertices are the characteristic vectors of stable sets of $T^{\circlearrowleft}$ containing the maximum number vertices in each edge gadget E_{ij}. Thus, F is defined by intersecting $\mathrm{STAB}(G)$ with the (face inducing) hyperplanes $\sum_{v \in V(E_{ij})} x_v = \ell + 2$ for all $ij \in E(T)$. Here x_v is the coordinate for vertex v in $T^{\circlearrowleft}$. For simplicity, we denote by $x_{ij}^{(k)}$ the coordinate for the additional vertex $u_{ij}^{(k)}$ of the 5-cycle in E_{ij}, see Figure 2.

Then it can be easily verified that F is an extension of corollary(T) via the affine map $\pi \colon x \mapsto y = \pi(x)$ where

$$y_{ij} = \begin{cases} x_i & \text{if } i = j, \\ 1 - x_{ij}^{(1)} - x_{ij}^{(2)} & \text{if } i \neq j. \end{cases}$$

In this definition, the y_{ij} are the correlation variables, with $i, j \in V(T)$ and either $i = j$ or $ij \in E(T)$.

Now Lemma 1 gives

$$\mathrm{xc}(\mathrm{STAB}(G)) \geqslant \mathrm{xc}(F) \geqslant \mathrm{xc}(\mathrm{corollary}(T)).$$

For $T = K_t$, using Theorem 2, we have

$$\mathrm{xc}(\mathrm{STAB}(G)) \geqslant \mathrm{xc}(\mathrm{corollary}(t)) \geqslant 2^{t \log(3/2)}.$$

◄

4.1   Existence of Gadgets in Random Graphs

In this section, we estimate the probability that a random Erdős–Rényi graph $G = G(n, p)$ contains an induced copy of a graph H. Recall that in the $G(n, p)$ model, each of the $\binom{n}{2}$ pairs of vertices is connected by an edge with probability p, independently from the other edges. The next lemma is key for proving lower bounds on the extension complexity of $\mathrm{STAB}(G(n, p))$ via embedding $H = T^{\circlearrowleft}$ as an induced subgraph. The lemma is formulated in a general for future applications to many types of subgraphs H.

▶ **Lemma 9.** *Let H be a graph with v vertices and with all induced subgraphs having average degree at most d. Let $0 < p \leqslant 1/2$ and*

$$g = g(n, p, v) := \frac{v^2 p^{-\frac{d}{2}}(1-p)^{-\frac{v}{2}}}{n - v}.$$

The probability of $G(n, p)$ not containing an induced copy of H satisfies

$$\mathbb{P}\left[H \overset{\text{ind}}{\nsubseteq} G(n, p)\right] \leqslant c_0 g^2 \approx 1.23 g^2,$$

where $c_0 := \exp(2W(1/\sqrt{2}))/2$ and W is the Lambert W-function, the inverse of $x \to x \exp x$.

Proof. The proof is via the second-moment method.

Let S be any graph isomorphic to H with $V(S) \subseteq V(G)$. Let X_S be the indicator random variable of S being an induced subgraph of G. Obviously, the total number X of induced subgraphs of G isomorphic to H satisfies $X = \sum_S X_S$. We estimate the expectation and variance of X. Let e denote the number of edges of H, and let $\text{Aut}(H)$ denote the automorphism group of H. The expectation is clearly

$$\mathbb{E}\left[X\right] = \sum_S \mathbb{E}\left[X_S\right] = \binom{n}{v} \frac{v!}{|\text{Aut}(H)|} p^e (1-p)^{\binom{v}{2} - e}.$$

The variance needs more preparations. Let now S and T be two graphs isomorphic to H with $V(S), V(T) \subseteq V(G)$. Using that X_S and X_T are independent and thus $\text{Cov}\left[X_S, X_T\right] = 0$ when $|V(S) \cap V(T)| \leqslant 1$ we get

$$\text{Var}\left[X\right] = \sum_{S,T} \text{Cov}\left[X_S, X_T\right] \leqslant \sum_{|V(S) \cap V(T)| \geqslant 2} \mathbb{E}\left[X_S X_T\right]$$

$$= \sum_{|V(S) \cap V(T)| \geqslant 2} \mathbb{E}\left[X_S\right] \mathbb{E}\left[X_T \mid X_S = 1\right] = \mathbb{E}\left[X\right] \sum_{T : \, |V(S) \cap V(T)| \geqslant 2} \mathbb{E}\left[X_T \mid X_S = 1\right].$$

Note that in the last sum S is fixed, and by symmetry, the sum is independent of the actual value of S. That is why we could factor it out. We obtain via Chebyshev's inequality,

$$\mathbb{P}\left[H \overset{\text{ind}}{\nsubseteq} G(n, p)\right] = \mathbb{P}\left[X = 0\right] \leqslant \frac{\text{Var}\left[X\right]}{\mathbb{E}\left[X\right]^2} \leqslant \frac{\sum_{T : \, |V(S) \cap V(T)| \geqslant 2} \mathbb{E}\left[X_T \mid X_S = 1\right]}{\mathbb{E}\left[X\right]}.$$

We shall estimate $\mathbb{E}\left[X_T \mid X_S = 1\right]$, which is the probability that H is induced in G provided S is induced in G, as a function of $k := |V(S) \cap V(T)|$. We assume that S and T coincide on $V(S) \cap V(T)$, and therefore have at most $dk/2$ edges in common, as their intersection is isomorphic to an induced subgraph of H, and therefore have average degree at most d by assumption. Hence as $p \leqslant 1/2$

$$\mathbb{E}\left[X_T \mid X_S = 1\right] = \mathbb{P}\left[T \overset{\text{ind}}{\subseteq} G \,\middle|\, S \overset{\text{ind}}{\subseteq} G\right] \leqslant p^{e - \frac{d}{2}k}(1-p)^{\binom{v}{2} - e - \binom{k}{2} + \frac{d}{2}k}.$$

This is clearly also true if S and T do not coincide on $V(S) \cap V(T)$, as then the probability

is 0. Now we can continue our estimation by summing up for all possible T with $k \geqslant 2$:

$$\frac{\sum_T \mathbb{E}\left[X_T \mid X_S = 1\right]}{\mathbb{E}\left[X\right]} \leqslant \frac{\sum_{k=2}^{v} \binom{v}{k}\binom{n-v}{v-k}\frac{v!}{|\text{Aut }H|}p^{e-\frac{d}{2}k}(1-p)^{\binom{v}{2}-e-\binom{k}{2}+\frac{d}{2}k}}{\binom{n}{v}\frac{v!}{|\text{Aut }H|}p^{e}(1-p)^{\binom{v}{2}-e}}$$

$$= \sum_{k=2}^{v} \frac{\binom{v}{k}\binom{n-v}{v-k}}{\binom{n}{v}}p^{-\frac{d}{2}k}\left(\underbrace{(1-p)^{\frac{d+1-k}{2}}}_{\leqslant(1-p)^{-\frac{v}{2}}}\right)^{k} \leqslant \sum_{k=2}^{v}\frac{v^k}{2(k-2)!}\left(\frac{v}{n-v}\right)^k\left(p^{-\frac{d}{2}}(1-p)^{-\frac{v}{2}}\right)^k$$

$$= \frac{1}{2}g^2\sum_{k=2}^{v}\frac{1}{(k-2)!}g^{k-2} \leqslant \frac{1}{2}g^2\exp(g),$$

as

$$\frac{\binom{v}{k}\binom{n-v}{v-k}}{\binom{n}{v}} \leqslant \frac{\binom{v}{k}\frac{(n-v)^{v-k}}{(v-k)!}}{\frac{(n-v)^v}{v!}} = \binom{v}{k}^2\frac{k!}{(n-v)^k} \leqslant \frac{1}{k!}\left(\frac{v}{n-v}\right)^k.$$

The lemma follows: the probability of H not being an induced subgraph is at most $e^g g^2/2$. This upper bound is 1 exactly if $g = 2W(1/\sqrt{2})$. For $g \leqslant 2W(1/\sqrt{2})$, we obtain the upper bound in the lemma. For $g \geqslant 2W(1/\sqrt{2})$, the upper bound in the lemma is at least 1, so the statement is obvious. ◀

4.2   High Extension Complexity with High Probability

In order to obtain lower bounds on the extension complexity of the stable set polytope of $G = G(n,p)$, we use Lemma 9 together with Lemma 8, taking H to be $K_t^{\bigcirc}$. We obtain the following result:

▶ **Main Theorem 10** (Super-polynomial xc of $\text{STAB}(G(n,p))$ w.h.p.)**.** *With high probability, the stable set polytope of the random graph $G(n,p)$ has at least the following extension complexity, depending on the size of p:*
(i) *For $p = \omega(1/\sqrt[4]{n})$ and fixed $0 < c < 2/\sqrt{3} \approx 1.1547$, we have*

$$\mathbb{P}\left[\text{xc}(\text{STAB}(G(n,p))) \geqslant 2^{\sqrt{c\frac{\ln(np^4)}{p}}\log(3/2)}\right] = 1 - o(1). \tag{1}$$

(ii) *For $c > 0$ and $c/\sqrt[3]{n} \leqslant p = o(1)$ we have*

$$\mathbb{P}\left[\text{xc}(\text{STAB}(G(n,p))) \geqslant 2^{\frac{\log(3/2)}{\sqrt{p\ln(1/p)}}}\right] = 1 - O(1/c^6). \tag{2}$$

(iii) *Moreover, for any fixed $c > 0$ for all $1/n < p \leqslant c/\sqrt[3]{n}$ and $0 < \delta < 1$*

$$\mathbb{P}\left[\text{xc}(\text{STAB}(G(n,p))) \geqslant 2^{\delta\sqrt{\frac{\sqrt{pn}}{\ln(1/p)}}\log(3/2)}\right] \geqslant 1 - O(\delta^8). \tag{3}$$

As an illustration of Main theorem 10, we include concrete lower bounds in special cases of interest.

▶ **Corollary 11.** *For every fixed $0 < \varepsilon < 1$, we have*

$$\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G(n, n^{-\varepsilon}))) \geqslant 2^{\sqrt{(1-4\varepsilon)n^\varepsilon \ln n}\, \log(3/2)}\right] = 1 - o(1) \qquad \textit{for } \varepsilon < 1/4, \qquad (4)$$

$$\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G(n, n^{-\varepsilon}))) \geqslant 2^{\frac{n^{\varepsilon/2}}{\sqrt{\varepsilon \ln n}}\, \log(3/2)}\right] = 1 - o(1) \qquad \textit{for } \varepsilon < 1/3, \qquad (5)$$

$$\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G(n, n^{-\varepsilon}))) \geqslant 2^{\frac{n^{(1-\varepsilon)/4}}{\ln n}\, \log(3/2)}\right] = 1 - o(1) \qquad \textit{for } \varepsilon \geqslant 1/3. \qquad (6)$$

Below the $p = n^{-\varepsilon}$ range, we obtain

$$\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G(n, (\ln^{6+\varepsilon} n)/n))) \geqslant 2^{\ln^{1+\varepsilon/5} n \cdot \log(3/2)}\right] = 1 - o(1), \qquad (7)$$

and (at the other end of the range) for fixed $\delta > 0$,

$$\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G(n, \delta \ln^{-1} n))) \geqslant n^{\delta^{-1/2} \log(3/2)}\right] = 1 - o(1). \qquad (8)$$

Proof. Equations (11) and (11) are special cases of (1). For Equation (11), we choose $p = n^{-\varepsilon}$ and $c = 1$. For Equation (1), we choose $p = \delta \ln^{-1} n$ and $c = 1.1$, a bit larger than 1, then the square root in (1) becomes

$$\sqrt{c\frac{\ln np^4}{p}} = \sqrt{c\frac{\ln n + 4\ln(\delta) + 4\ln \ln^{-1} n}{\delta}} \ln n = c\delta^{-1/2}(1 + o(1))\ln n > \delta^{-1/2} \ln n,$$

proving the equation.

Equation (11) follows from Equation (2) via $p = n^{-\varepsilon}$. Equations (11) and (11) are special cases of Equation (3). Equation (11) is the case $p = n^{-1+\varepsilon}$ and $\delta = \sqrt{\varepsilon} \ln^{-1} n$. For Equation (11), we choose $p = (\ln^{6+\varepsilon} n)/n$ and $\delta = \ln^{-\varepsilon/20} n$, then the interesting part of the exponent is

$$\delta\sqrt{\frac{\sqrt{pn}}{\ln(1/p)}} = \ln^{-\varepsilon/20} n\sqrt{\frac{\sqrt{\ln^{6+\varepsilon} n}}{\ln n - \ln\ln^{6+\varepsilon} n}} > \ln^{-\varepsilon/20} n\sqrt{\frac{\sqrt{\ln^{6+\varepsilon} n}}{\ln n}} = \ln^{1+\varepsilon/5} n$$

proving the claim. ◀

Now we are going to prove the main theorem of Section 4.2.

Proof of Main Theorem 10. We apply Lemma 9 to the graph $H := K_t^\bigcirc$ together with Lemma 8 to obtain:

$$\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G(n,p))) \geqslant 2^{t\log(3/2)}\right] \geqslant \mathbb{P}\left[K_t^\bigcirc \overset{\mathrm{ind}}{\subseteq} G(n,p)\right]$$
$$\geqslant 1 - c_0 \frac{v^4 p^{-d}(1-p)^{-v}}{(n-v)^2}$$
$$\geqslant 1 - c_0(1 + o(1))\frac{v^4 p^{-d}e^{pv}}{n^2} \qquad \text{if } v = o(n).$$

Here v is the number of vertices of H, and every induced subgraph of H should have average degree at most d. We shall estimate the last fraction $v^4 p^{-d}e^{pv}/n^2$, using the d provided by Lemma 7. Below we will tacitly assume $t = \omega(1)$, which is w.l.o.g because $\mathrm{xc}(\mathrm{STAB}(G(n,p))) \geqslant n$ always.

Now we shall substitute various values for p, t, d, ℓ to obtain the equations of the theorem. We will verify $v = o(n)$ and $v^4 p^{-d} e^{pv}/n^2 = o(1)$ to obtain an $1 - o(1)$ lower bound from the last inequality.

For establishing (1), we choose

$$\ell := 0 \qquad\qquad t := \left\lceil c\sqrt{\frac{\ln(np^4)}{p}}\right\rceil \qquad\qquad d := 4.$$

Note that for $p \geqslant 1/\sqrt[4]{n}$,

$$v = t + 3\binom{t}{2} = \left(\frac{3}{2} + o(1)\right)t^2 = \left(\frac{3}{2} + o(1)\right)c^2 \frac{\ln(np^4)}{p} \leqslant \left(\frac{3}{2} + o(1)\right)c^2 \sqrt[4]{n}\ln n = o(n),$$

and hence

$$\frac{v^4 p^{-d} e^{pv}}{n^2} = \left(\frac{3}{2} + o(1)\right)^4 (pt^2)^4 e^{(3/2+o(1))pt^2 - 2\ln(np^4)}$$

$$\leqslant \left(\frac{3}{2} + o(1)\right)^4 c^8 \left(\ln(np^4)\right)^4 \exp\left\{\left[\left(\frac{3}{2} + o(1)\right)c^2 - 2\right]\ln(np^4)\right\} = o(1),$$

as $np^4 = \omega(1)$ by assumption. This finishes the proof of (1).

We turn to (2) and (3). We will choose a positive ℓ to approximately minimize the fraction in terms of the other parameters. To ease computation, let

$$\gamma := \frac{2\ell + 3}{2} > 1.$$

Then the parameters v and d look like

$$d = 2 + \frac{4}{2\ell + 3} = 2 + \frac{2}{\gamma},$$

$$v = t + (2\ell + 3)\binom{t}{2} = \gamma t^2 + (1 - \gamma)t < \gamma t^2.$$

Hence

$$\frac{v^4 p^{-d} e^{pv}}{n^2} < \frac{\gamma^4 t^8 e^{p\gamma t^2 + 2(\ln(1/p))/\gamma}}{p^2 n^2}.$$

The γ minimizing the expression is

$$\frac{\sqrt{4 + 2pt^2\ln(1/p)} - 2}{pt^2} = \frac{2\ln(1/p)}{\sqrt{4 + 2pt^2\ln(1/p)} + 2},$$

but we use an approximation as ℓ needs to be an even integer. Therefore we choose

$$\ell = 2\left\lceil \frac{\ln(1/p)}{\sqrt{4 + 2pt^2\ln(1/p)} + 2} - \frac{3}{4}\right\rceil.$$

We will verify later that actually $\ell = \omega(1)$. Hence

$$\gamma = (1 + o(1))\frac{2\ln(1/p)}{\sqrt{4 + 2pt^2\ln(1/p)} + 2} = (1 + o(1))\frac{\sqrt{4 + 2pt^2\ln(1/p)} - 2}{pt^2},$$

and

$$\frac{v^4 p^{-d} \mathrm{e}^{pv}}{n^2} < (1 + o(1)) \left(\frac{2pt^2 \ln(1/p)}{\sqrt{np^3}} \right)^4 \frac{\mathrm{e}^{(2+o(1))\sqrt{4 + 2pt^2 \ln(1/p)}}}{\left(\sqrt{4 + 2pt^2 \ln(1/p)} + 2 \right)^4}$$

$$= (1 + o(1)) \frac{\mathrm{e}^{(2+o(1))\sqrt{4 + 2pt^2 \ln(1/p)}} \left(\sqrt{4 + 2pt^2 \ln(1/p)} - 2 \right)^4}{(np^3)^2}. \tag{9}$$

Now we shall substitute various values for p and t to obtain the equations of the theorem. We will need to verify $\ell = \omega(1)$ and $v = o(n)$ for every choice.

For Equation (3), i.e., in the case $1/n < p \leqslant c/\sqrt[3]{n}$, we neglect the exponential term in (4.2) for the choice of t:

$$t = \left\lceil \delta \sqrt{\frac{\sqrt{pn}}{\ln(1/p)}} \right\rceil.$$

Here $0 < \delta < 1$ is an additional parameter. Rearranging gives us

$$2pt^2 \ln(1/p) = (1 + o(1))\delta^2 \sqrt{np^3} \leqslant (1 + o(1))\delta^2 c^{3/2} \leqslant (1 + o(1))c^{3/2},$$

so in particular,

$$\ell \geqslant 2 \left\lceil \frac{\ln(\sqrt[3]{n}/c)}{\sqrt{4 + (1 + o(1))c^{3/2}} + 2} - \frac{3}{4} \right\rceil = \omega(1)$$

$$v < \gamma t^2 = O(1/p) = O(\sqrt[3]{n}) = o(n).$$

Finally,

$$\frac{v^4 p^{-d} \mathrm{e}^{pv}}{n^2} < (1 + o(1)) \frac{\mathrm{e}^{(2+o(1))\sqrt{4 + (1+o(1))c^{3/2}}} \left(\sqrt{4 + (1 + o(1))\delta^2 \sqrt{np^3}} - 2 \right)^4}{(np^3)^2}$$

$$\leqslant (1 + o(1))\mathrm{e}^{(2+o(1))\sqrt{4 + (1+o(1))c^{3/2}}} \left((1/4 + o(1))\delta^2 \right)^4 = O(\delta^8),$$

as claimed.

For Equation (2), i.e., when $c/\sqrt[3]{n} \leqslant p = o(1)$, we choose

$$t = \left\lceil \frac{1}{\sqrt{p \ln(1/p)}} \right\rceil.$$

This provides the estimate

$$2pt^2 \ln(1/p) = 2 + o(1),$$

hence $\ell = \Theta(\ln(1/p)) = \omega(1)$, and $v < \gamma t^2 = O(1/p) = O(\sqrt[3]{n}) = o(n)$. Finally,

$$\frac{v^4 p^{-d} \mathrm{e}^{pv}}{n^2} = (1 + o(1)) \frac{\mathrm{e}^{(2+o(1))\sqrt{4 + (2+o(1))}} \left(\sqrt{4 + (2 + o(1))} - 2 \right)^4}{(np^3)^2}$$

$$= O \left(\frac{1}{(np^3)^2} \right) = O(1/c^6),$$

as $np^3 \geqslant c^3$. ◀

Main Theorem 10 gives super-polynomial lower bounds all the way from $p = \Omega\left(\frac{\log^{6+\varepsilon} n}{n}\right)$ to $p = O\left(\frac{1}{\log n}\right)$. The key for being able to cover the whole regime is to have the gadgets depend on the parameter choice. Notice that for $p < 1/n$ a random graph almost surely will have all its components of size $O(\log n)$, making the stable set problem easy to solve, so that we essentially leave only a small polylog gap.

4.3   Upper Bound on Extension Complexity with High Probability

We now complement Main Theorem 10 with an upper bound, which is close to the lower bound, up to an essentially quadratic gap in the exponent.

▶ **Theorem 12** (Upper bound on the xc of $\mathrm{STAB}(G(n,p))$ w.h.p.). *For* $0 < p \leqslant 1/2$,

$$\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G)) \geqslant 2^{\Omega\left(\frac{\ln^2 n}{p}\right)}\right] \leqslant n^{-\Omega\left(\frac{\ln n}{p}\right)}.$$

In particular, for $p = n^{-\varepsilon}$, *we obtain* $\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G)) \geqslant 2^{\Omega\left(n^{\varepsilon} \ln^2 n\right)}\right] = o(1)$ *and similarly for* $p = \delta \ln^{-1} n$, *we get* $\mathbb{P}\left[\mathrm{xc}(\mathrm{STAB}(G)) \geqslant n^{\Omega\left(\frac{\ln^3 n}{\delta}\right)}\right] = o(1)$.

The upper bound stated in Theorem 12 essentially relies on the following basic result.

▶ **Lemma 13.** *Every polytope P has an extension complexity at most the number of its vertices.*

Proof. Let V be the set of vertices of P, and let Q be a simplex with $|V|$ vertices. The simplex Q is an extension of P via mapping the vertices of Q one-to-one to V in an arbitrary fashion, and extending to an affine mapping on Q. This extension has size $|V|$.   ◀

We are ready to prove our theorem.

Proof of Theorem 12. By standard arguments (see, e.g., [8, Chapter 11, page 300]), for $G = G(n,p)$ we have

$$\mathbb{P}\left[\alpha(G) \geqslant r\right] \leqslant \left(n\, e^{-p(r-1)/2}\right)^r$$

and thus for $r = 4\frac{\ln n}{p}$ we get

$$\mathbb{P}\left[\alpha(G) \geqslant 4\frac{\ln n}{p}\right] \leqslant \left(\frac{n}{\sqrt{e}}\right)^{-4\frac{\ln n}{p}}.$$

Therefore, with very high probability, we have $\alpha(G) \leqslant 4\frac{\ln n}{p}$. Using the inequality $\sum_{i=0}^{k} \binom{n}{i} \leqslant (n+1)^k$, we get

$$\#(\text{stable sets in } G) \leqslant (n+1)^{\alpha(G)} = 2^{\log(n+1)\alpha(G)} = 2^{\left(\frac{1}{\ln 2} + o(1)\right)\ln(n)\alpha(G)}.$$

The result then follows directly from Lemma 13.   ◀

5 Concluding Remarks

We conclude with the following conjecture whose validity, we believe, is necessary to strengthen the result, close the remaining gap, as well as establishing truly exponential lower bounds on the extension complexity of further combinatorial problems.

▶ **Conjecture 14** (Sparse Graph Conjecture). *There exists an infinite family* $(T_k)_{k \in \mathbb{N}}$ *of template graphs such that, denoting by* t_k *the number of vertices of* T_k*: (i)* $\mathrm{xc}(\mathrm{corollary}(T_k)) = 2^{\Omega(t_k)}$*; (ii)* T_k *has bounded average degree; (iii)* $t_k \leqslant t_{k+1}$ *but at the same time* $t_{k+1} = O(t_k)$.

The existence of such a family would have various consequences.

Exact Case

Assuming the Sparse Graph Conjecture we would obtain that the extension complexity of polytopes for important combinatorial problems considered in [9, 1, 13] including (among others) the stable set polytope, knapsack polytope, and the 3SAT polytope would have truly exponential extension complexity, that is $2^{\Omega(n)}$ extension complexity, where n is the *dimension* of the polytope.

The recent groundbreaking result of [15] gives $2^{\Omega(n)}$ bounds for the extension complexity of the matching polytope and TSP polytope. These bounds are also tight up to constants, but this time the upper bound does not come from the number of vertices but rather from the number of facets and dynamic programming algorithms, respectively. Notice that the dimension of both polytopes is $d = \Theta(n^2)$, thus the bounds are in fact $2^{\Omega(\sqrt{d})}$.

Average Case

As observed above, there is a quadratic gap in the best current lower and upper bounds on the worst-case extension complexity of the stable set polytope: $2^{\Omega(\sqrt{n})}$ versus 2^n respectively. This is reflected in the results we obtain here. Assuming the Sparse Graph Conjecture we could reduce the gap between upper and lower bounds to a logarithmic factor. Moreover, our results could be strengthened to establish super-polynomial lower bounds on the average-case extension complexity up to constant probability p.

—— **References** ——

1   D. Avis and H. R. Tiwary. On the extension complexity of combinatorial polytopes. *ArXiv e-prints*, February 2013.
2   G. Braun, S. Fiorini, S. Pokutta, and D. Steurer. Approximation Limits of Linear Programs (Beyond Hierarchies). In *53rd IEEE Symp. on Foundations of Computer Science (FOCS 2012)*, pages 480–489, 2012.
3   G. Braun and S. Pokutta. Common information and unique disjointness. In *IEEE 54th Annual Symp. on Foundations of Computer Science (FOCS 2013)*, pages 688–697, 2013. http://eccc.hpi-web.de/report/2013/056/.
4   G. Braun and S. Pokutta. The matching polytope does not admit fully-polynomial size relaxation schemes. *preprint available at http://arxiv.org/abs/1403.6710*, 2014.
5   M. Braverman and A. Moitra. An information complexity approach to extended formulations. In *Proceedings of the 45th annual ACM symposium on Theory of computing*, pages 161–170, 2013.

6   S. O. Chan, J. R. Lee, P. Raghavendra, and D. Steurer. Approximate constraint satisfaction requires large LP relaxations. In *IEEE 54th Annual Symp. on Foundations of Computer Science (FOCS 2013)*, pages 350–359, 2013.

7   M. Conforti, G. Cornuéjols, and G. Zambelli. Extended formulations in combinatorial optimization. *4OR*, 8:1–48, 2010.

8   R. Diestel. *Graph Theory*. Springer-Verlag Heidelberg, New York, 2005.

9   S. Fiorini, S. Massar, S. Pokutta, H. R. Tiwary, and R. de Wolf. Linear vs. Semidefinite Extended Formulations: Exponential Separation and Strong Lower Bounds. *Proc. STOC 2012*, 2012.

10 M. X. Goemans and D. P. Williamson. Improved approximation algorithms for maximum cut and satisfiability problems using semidefinite programming. *J. Assoc. Comput. Mach.*, 42:1115–1145, 1995.

11 V. Kaibel. Extended formulations in combinatorial optimization. *Optima*, 85:2–7, 2011.

12 V. Kaibel and S. Weltge. A Short Proof that the Extension Complexity of the Correlation Polytope Grows Exponentially. *ArXiv e-prints*, July 2013.

13 S. Pokutta and M. Van Vyve. A note on the extension complexity of the knapsack polytope. *Operations Research Letters*, 41:347–350, 2013.

14 T. Rothvoß. Some 0/1 polytopes need exponential size extended formulations, 2011. arXiv:1105.0036.

15 Thomas Rothvoß. The matching polytope has exponential extension complexity. *ArXiv e-prints*, 2013.

16 J. C. Williams. A linear-size zero-one programming model for the minimum spanning tree problem in planar graphs. *Networks*, 39:53–60, 2002.