Noname manuscript No.
(will be inserted by the editor)

A Semidefinite Hierarchy for Expected Independence Numbers of Random Graphs

Kevin Shu · Diego Cifuentes · Alejandro Toriello

the date of receipt and acceptance should be inserted later

Abstract We introduce convex optimization methods to find upper bounds on the expected independence number of a random graph, in the vein of the Lovász theta function's bound for the independence number of a deterministic graph. Specifically, we propose a hierarchy of semidefinite programs whose values upper bound the expected independence number. Our hierarchy can be applied to arbitrary random graph models, and only requires bounds on the probabilities that subsets of vertices are independent in the resulting graph. For symmetric random graphs, the last level of the hierarchy is equivalent to a linear program whose optimal value can often be calculated or approximated in closed form. We show that our methods provide good upper bounds in a number of examples, including Erdős-Rényi graphs and geometric random graphs.

Keywords independence number · random graphs · semidefinite programming · Lovász theta function.

Mathematics Subject Classification (2020) 65K05, 90C22, 05C80

1 Introduction

For a simple graph G, an independent set (also called a stable set or a node packing) is a subset of vertices with no edges between any two vertices in the subset. The independence number of G, denoted $\alpha(G)$, is the cardinality of the largest independent set in G. In this paper we are interested in the independence number of random graphs given by a variety of models. More precisely, we seek upper bounds on the expected independence number of a random graph that are analogous to and build on the Lovász theta function bound for the independence number of a deterministic graph.

Georgia Institute of Technology, Atlanta GA 30318, USA E-mail: kshu8@gatech.edu, {diego.cifuentes, atoriello}@isye.gatech.edu

The Lovász theta number $\vartheta(G)$ of a graph G with vertex set $[n] := \{1, \dots, n\}$ and edge set E is the optimal value of the semidefinite program (SDP):

$$\begin{aligned}
\vartheta(G) \quad = \quad \max_{Z \in \mathbb{S}^n, z \in \mathbb{R}^n} \quad & \operatorname{tr}(Z) \\
\text{s.t.} \quad & Z_{ij} = 0, \quad \forall\, ij \in E \\
& Z_{ii} = z_i, \quad \forall\, i \in [n] \\
& \begin{pmatrix} 1 & z^\intercal \\ z & Z \end{pmatrix} \succeq 0.
\end{aligned}$$

The theta number was introduced in [13] in a different form; the SDP formulation above is from [7]. For certain graph classes, notably perfect graphs, $\vartheta(G)$ provides the exact value of $\alpha(G)$ [6].

If G is chosen randomly according to some distribution, the independence number $\alpha(G)$ is a random variable. The expected independence number has been the subject of extensive study [1,2], and serves as a dual bound in related dynamic graph optimization problems [16]. However, many existing results depend on ad hoc techniques tailored to specific probability models; furthermore, the expected value of $\vartheta(G)$ and similar SDP bounds for deterministic graphs may be far from the expected independence number [5, 9]. With this motivation, we take a non-parametric approach that can be applied to arbitrary random graph models. We introduce an SDP hierarchy that only depends on certain marginal probabilities, but nevertheless yields optimal or near-optimal results for several notable random graph classes. This bound sequence is inspired by the Lasserre hierachy for Boolean quadratic programs, which has also been intently studied for the independent set [12, 14] and similar problems [19].

Suppose G is a random graph with vertex set $[n]$ drawn from a probability distribution $\mathcal{G}$; that is, the edges of G are a random subset of $\{(i, j) : i, j \in [n]\}$, and the appearance of the edges need not be independent. We let p_{ij} denote the probability that the edge (i, j) appears in G. We show in Section 2 that the following variant of the Lovász theta function provides an upper bound on the expected independence number of G drawn from the distribution $\mathcal{G}$:

$$\begin{aligned}
\phi_1(\mathcal{G}) \quad = \quad \max_{Z \in \mathbb{S}^n, z \in \mathbb{R}^n} \quad & \operatorname{tr}(Z) \\
\text{s.t.} \quad & 0 \le Z_{ij} \le 1 - p_{ij}, \quad \forall\, ij \in \tbinom{[n]}{2} \\
& Z_{ii} = z_i, \quad \forall\, i \in [n] \\
& \begin{pmatrix} 1 & z^\intercal \\ z & Z \end{pmatrix} \succeq 0.
\end{aligned} \tag{1}$$

If $p_{ij} \in \{0, 1\}$ for all i, j, this leads to the variant of the deterministic theta function which includes nonnegativity constrainst on the entries of Z. Our proposed bound hierarchy builds on ϕ_1, yielding progressively tighter bounds on the expected independence number.

Assume now that the random graph is symmetric, i.e. the probability distribution on the edges is invariant under a permutation of the vertices. Formally, we have that for any fixed graph G and any permutation $\pi \in S_n$,

$$\Pr_{G \sim \mathcal{G}}(G = T) = \Pr_{G \sim \mathcal{G}}(G = \pi(T)),$$

where $\pi(T)$ denotes the graph T with π applied to the vertices. In this case, the final level of our hierarchy reduces to the linear program (LP)

$$\psi(\mathcal{G}) \quad = \quad \max_{x \in \Delta^n} \quad \sum_{i=1}^{n} i \cdot x_i \tag{2a}$$

$$\text{s.t.} \quad \sum_{i=\ell}^{n} \frac{\binom{n-\ell}{i-\ell}}{\binom{n}{i}} x_i \le q_\ell, \quad \forall\ \ell \in [n], \tag{2b}$$

where $\Delta^n := \left\{x \in \mathbb{R}^n : \sum_i x_i = 1,\ x \ge 0\right\}$ is the probability simplex and q_ℓ is the probability that a subset of size ℓ is independent in G. Intuitively, the variable x_i in this LP represents the probability that $\alpha(G)$ is exactly equal to i, and the last constraint represents the probability that $\{1, \ldots, \ell\}$ is contained in a randomly chosen maximum independent set of G (in which case, it is clearly independent).

We show that this LP captures a number of combinatorial inequalities, and examine how it behaves in the particular cases of Erdős-Rényi random graphs, uniformly random spanning trees and geometric random graphs in Section 4. For example, for Erdős-Rényi random graphs, our LP provides an asymptotically tight bound, whereas the expected value of the theta function provides a much weaker bound [9].

2 Semidefinite Hierarchy for General Random Graphs

Here, we introduce the hierarchy that we explore in this paper. Let G be a random graph with vertex set $[n]$, and fix some $t \in [n]$, which represents the level of the hierarchy. For each subset $S \subseteq [n]$, we let

$$q_S := \Pr_{G \sim \mathcal{G}}(S \text{ is independent in } G).$$

We then introduce a variable y_S for each $S \subseteq [n]$ with $|S| \le 2t$. If S is the singleton set $\{i\}$, then we let $y_i = y_S$ for shorthand.

Let $m = \sum_{s=0}^{t} \binom{n}{s}$ be the number of subsets of $[n]$ of size at most t. Let $M_t(y)$ be the $m \times m$ matrix whose rows and columns are indexed by subsets of $[n]$ of size at most t, so that $M_t(y)_{A,B} = y_{A \cup B}$. This matrix is sometimes referred to as the *moment matrix* for the binary vector y. We use the definition given in [18], which originates from [10, 17]. The idea of using a semidefinite programming hierachy for the independent set problem was first described in [14], and the connection to Lasserre's hierarchy in the context of polynomial optimization was given in [11]. Denote by $v_T = (1_{S \subseteq T})_{S \subseteq [n]} \in \mathbb{R}^{2^n}$ the indicator vector for each $T \subseteq [n]$. The important property of this matrix is that if $y \in \mathbb{R}^{2^n}$ and $y_\emptyset = 1$, then

$$\mu_n \ := \ \{y \in \mathbb{R}^{2^n} : y_\varnothing = 1,\ M_n(y) \succeq 0\} \ = \ \operatorname{conv}(v_T : T \subseteq [n]). \tag{3}$$

This is known as the *moment polytope* [18]. The above equation implies that

$$M_n(y) \succeq 0 \quad \iff \quad \exists \text{ distribution } \pi \text{ on } 2^{[n]} \text{ s.t. } y_S = \Pr_{T \sim \pi}(S \subseteq T). \tag{4}$$

We consider the following SDP, which forms the t-th level of our hierarchy:

$$\begin{aligned} \phi_t(\mathcal{G}) \quad := \quad \max_y \quad & \sum_{i=1}^n y_i \\ \text{s.t.} \quad & y_\varnothing = 1 \\ & 0 \le y_S \le q_S, \quad \forall\, S \subseteq [n] \text{ with } |S| \le 2t. \\ & M_t(y) \succeq 0. \end{aligned} \tag{5}$$

Note that when $t = 1$, we obtain (1).

Theorem 1 *The hierarchy defined above satisfies*

$$\mathbb{E}[\alpha(G)] \le \phi_n(\mathcal{G}) \le \cdots \le \phi_2(\mathcal{G}) \le \phi_1(\mathcal{G}).$$

Proof If $\ell \ge m$, $\phi_\ell(\mathcal{G}) \le \phi_m(\mathcal{G})$, since there are more constraints on the variables y_i in $\phi_\ell(\mathcal{G})$ compared to $\phi_m(\mathcal{G})$. So, it suffices to show $\mathbb{E}[\alpha(G)] \le \phi_n(\mathcal{G})$. We define a random subset $T \subseteq [n]$ as follows: we first sample an instance of the random graph G and then choose a uniformly random maximum independent set T of G. Define $\hat{y}_S$ to be the probability that the random subset T contains S as a subset. We claim that $\hat{y}$ is feasible for the program $\phi_n(\mathcal{G})$, and that its objective value is precisely $\mathbb{E}[\alpha(G)]$. From this, it follows that $\mathbb{E}[\alpha(G)] \le \phi_n(\mathcal{G})$.

We clearly have that $\hat{y}_\varnothing = 1$, and for any $S \subseteq [n]$, $\hat{y}_S \ge 0$. To see that $q_S \ge \hat{y}_S$, if S is contained in a maximum independent set, then S is independent, and thus

$$q_S = \Pr(S \text{ is independent}) \ge \Pr(S \subseteq T) = \hat{y}_S.$$

The fact that $M_n(y) \succeq 0$ follows from (4). Hence, $\hat{y}$ is feasible for the SDP.

Finally, y_i is precisely the probability that $i \in T$, so

$$\mathbb{E}[\alpha(G)] = \mathbb{E}\Bigg[\sum_{i\in[n]} 1_{i\in T}\Bigg] = \sum_{i\in[n]} y_i. \qquad \square$$

The bound $\phi_n(\mathcal{G})$ corresponding to the last level of the hierarchy does not necessarily agree with the expected independence number $\mathbb{E}_{G\sim\mathcal{G}}[\alpha(G)]$. Nonetheless, $\phi_n(\mathcal{G})$ is tight with respect to a different random graph $\mathcal{G}'$ constrained by the same marginal probabilities q_S. Indeed, if y is optimal for ϕ_n, there is a probability distribution π on $2^{[n]}$ with $y_S = \Pr_{T\sim\pi}(S \subseteq T)$. Consider the following random model $\mathcal{G}'$: sample T from π and let G be the complete graph on n vertices with edges between vertices in T removed. Then $\mathbb{E}_{G\sim\mathcal{G}'}[\alpha(G)] = \phi_n(\mathcal{G})$, and for any $S \subseteq [n]$, $\Pr_{G\sim\mathcal{G}'}(S \text{ is independent in } G) = y_S \le q_S$.

The bounds $\phi_t(\mathcal{G})$ require the probabilities q_S, which may be difficult to compute in certain cases. However, any known upper bound for q_S also yields a bound on $\mathbb{E}_{G\sim\mathcal{G}}[\alpha(G)]$.

3 Linear Program for Symmetric Random Graphs

A symmetric random graph is one in which the probability of a graph appearing is invariant under permutations of its vertices. An example is the Erdős-Rényi random graph $G_{n,p}$, defined as a graph on n vertices in which each edge appears independently with probability p. We show that for symmetric graphs, the last level of our SDP hierarchy is equivalent to an LP.

Theorem 2 *Let G be a symmetric random graph on n vertices with probability distribution $\mathcal{G}$. Then $\phi_n(\mathcal{G}) = \psi(\mathcal{G})$, the optimal value of the LP in (2).*

Proof Let $\mu_n \subseteq \mathbb{R}^{2^n}, v_T \in \mathbb{R}^{2^n}$ as in (3). We first show that $\phi_n(\mathcal{G}) \geq \psi(\mathcal{G})$. Let $x \in \Delta^n$ be feasible for the LP; consider

$$y = \sum_{i=1}^{n} x_i \left(\frac{1}{\binom{n}{i}} \sum_{T \subseteq [n], |T|=i} v_T \right) \in \mathbb{R}^{2^n}.$$

Then $y \in \mu_n$, as it is a convex combination of the v_T's. Moreover, for any subset $S \subseteq [n]$ of size ℓ, there are precisely $\binom{n-\ell}{i-\ell}$ subsets of size i that contain S, so that

$$y_S = \sum_{i=\ell}^{n} \frac{\binom{n-\ell}{i-\ell}}{\binom{n}{i}} x_i \leq q_\ell = q_S,$$

where the last equality uses the fact that the distribution is symmetric. Hence, y is feasible for the SDP of $\phi_n(\mathcal{G})$. The objective value of this y is

$$\sum_{i=1}^{n} y_i = n y_1 = n \sum_{i=1}^{n} \frac{\binom{n-1}{i-1}}{\binom{n}{i}} x_i = \sum_{i=1}^{n} i x_i.$$

Since any feasible solution of the LP can be mapped into a feasible solution of the SDP with the same objective value, then $\psi(\mathcal{G}) \leq \phi_n(\mathcal{G})$.

We proceed to show that $\psi(\mathcal{G}) \geq \phi_n(\mathcal{G})$. Let $y \in \mathbb{R}^{2^n}$ be feasible for the SDP of $\phi_n(\mathcal{G})$. We have that $y \in \mu_n$, which implies that $y = \sum_{S \subseteq [n]} \lambda_S v_S$, for some λ in the probability simplex Δ^{2^n}. Let $x \in \Delta^n$ be such that x_i is the sum of all λ_S with $|S| = i$. We claim that x is feasible for the LP.

For a permutation $\pi \in S_n$, we let $\pi y \in \mathbb{R}^{2^n}$ be defined by setting $(\pi y)_S = y_{\pi S}$, where the permutation simply permutes the elements of each set S. Because G is symmetric, for any $\pi \in S_n$, πy is feasible for the SDP. Therefore, the following average is also feasible for the SDP:

$$\frac{1}{n!} \sum_{\pi \in S_n} \pi y = \sum_{i=1}^{n} x_i \left(\sum_{S \subseteq [n], |S|=i} v_S \right).$$

An argument similar to the first part of the proof shows that x is indeed feasible for the LP, and has the same objective value for the Lk as y, concluding the proof. □

We next prove a family of bounds on the optimal value of (2), including a particularly simple bound that we use in the next section. Afterwards, we show that (2) has an optimal solution with at most two nonzero entries.

Theorem 3 *Let G be a symmetric random graph on n vertices with probability distribution $\mathcal{G}$. Fix $i \leq n$ and $\ell \leq \min\{i+1, n\}$; then*

$$\psi(\mathcal{G}) \leq i - \frac{\binom{i}{\ell}}{\binom{i}{\ell-1}} + \frac{\binom{n}{\ell}}{\binom{i}{\ell-1}} q_\ell.$$

In particular, for any $i < n$, $\psi(\mathcal{G}) \leq i + \binom{n}{i+1} q_{i+1}$.

Proof The second bound follows from considering the case in which $\ell = i+1$, so we only need to consider the first. Consider the following two constraints in (2):

$$c_1(x) \coloneqq \sum_{j=1}^{n} x_j = 1, \qquad c_2(x) \coloneqq \sum_{j=\ell}^{n} \frac{\binom{n-\ell}{j-\ell}}{\binom{n}{j}} x_j \leq q_\ell.$$

It will be more convenient to work with the inequality

$$c_2'(x) \coloneqq \sum_{j=\ell}^{n} \binom{j}{\ell} x_j \leq \binom{n}{\ell} q_\ell,$$

which is equivalent to the inequality $c_2(x) \leq q_\ell$ because $\frac{\binom{n-\ell}{j-\ell}\binom{n}{\ell}}{\binom{n}{j}} = \binom{i}{\ell}$. This algebraic identity holds because $\binom{n-\ell}{i-\ell}\binom{n}{\ell} = \binom{i}{\ell}\binom{n}{i}$ are both the number of pairs (S,T) where $S \subseteq T \subseteq [n]$ with $|S| = \ell$ and $|T| = i$.

Combining these inequalities with coefficients $i - \frac{\binom{i}{\ell}}{\binom{i}{\ell-1}}$ and $\frac{1}{\binom{i}{\ell-1}}$ yields

$$t(x) \coloneqq \left(i - \frac{\binom{i}{\ell}}{\binom{i}{\ell-1}}\right) c_1(x) + \frac{1}{\binom{i}{\ell-1}} c_2'(x) \leq i - \frac{\binom{i}{\ell}}{\binom{i}{\ell-1}} + \frac{\binom{n}{\ell}}{\binom{i}{\ell-1}} q_\ell.$$

To show the theorem, it suffices to argue that $t(x)$ is an upper bound for the objective value, i.e. that the coefficient of x_j in $t(x)$ is at least j for $j \leq n$.

For all $j \leq \ell - 1$, the coefficient of x_j in $t(x)$ is $i - \frac{\binom{i}{\ell}}{\binom{i}{\ell-1}}$. This coefficient is

$$i - \frac{\binom{i}{\ell}}{\binom{i}{\ell-1}} = i - \frac{i-\ell+1}{\ell} = \frac{i+1}{\ell}(\ell-1) \geq \ell - 1 \geq j.$$

For $j \geq \ell$, the coefficient of x_j in $t(x)$ is $i - \binom{i}{\ell}\binom{i}{\ell-1}^{-1} + \binom{j}{\ell}\binom{i}{\ell-1}^{-1}$. This is at least j if and only if $(j-i)\binom{i}{\ell-1} \leq \binom{j}{\ell} - \binom{i}{\ell}$. When $j \geq i$, there are $(j-i)\binom{i}{\ell-1}$ subsets of $[j]$ of size ℓ of the form $S \cup \{b\}$, where S is a subset of $[i]$ of size $\ell - 1$ and $b \in \{i+1, \ldots, j\}$, and none of these are subsets of $[i]$.

When $j \leq i$, every subset of $[i]$ of size ℓ that is not in $[j]$ contains an element of $\{j+1, \ldots, i\}$; there are $(i-j)$ ways of choosing this element and $\binom{i}{\ell-1}$ ways of choosing the remaining elements. ☐

Theorem 4 *Let G be a symmetric random graph. There is an optimal solution to the LP* (2) *with at most two nonzero entries.*

Proof We construct the desired x explicitly. To begin, let

$$i^* = \max\left\{ i : \forall \ell \leq i, \binom{n}{\ell} q_\ell \geq \binom{i}{\ell} \right\}.$$

The maximum is over a non-empty set, because when $i = 1$, $q_\ell = 1$, so $\binom{n}{1} q_1 \geq \binom{1}{1}$.

If $i^* = n$, $q_n \geq 1$, so G is a clique on n vertices with probability 1. In this case, the coordinate vector $x = e_n$ is an optimal solution to the LP. Now suppose $i^* < n$. For $\ell \leq i^* + 1$, let

$$w_\ell = \frac{\binom{i^*+1}{\ell} - \binom{n}{\ell} q_\ell}{\binom{i^*+1}{\ell} - \binom{i^*}{\ell}}.$$

Define $x \in \mathbb{R}^n$ such that $x_i = 0$ for $i \notin \{i^*, i^*+1\}$, $x_{i^*} + x_{i^*} = 1$, and

$$x_{i^*} = \max\{w_\ell : \ell \leq i^* + 1\}.$$

We let ℓ^* satisfy $x_{i^*} = w_{\ell^*}$. We now claim that x is an optimal solution to (2). We first need to show that $x_{i^*} \in [0,1]$, in which case $x \in \Delta^n$. To see that $x_{i^*} \geq 0$, note that the definition of i^* implies that there must be some $\ell \leq i^* + 1$ so that $\binom{n}{\ell} q_\ell < \binom{i^*+1}{\ell}$, and for this value of ℓ, $w_\ell \geq 0$. To see that $x_{i^*} \leq 1$, note that for $\ell \leq i^*$, $\binom{n}{\ell} q_\ell \geq \binom{i^*}{\ell}$. This implies that $w_\ell \leq 1$ for $\ell \leq i^*$. On the other hand, for $\ell = i^* + 1$, $w_\ell = 1 - \binom{n}{i^*+1} q_{i^*+1} \leq 1$. This implies that $x_{i^*} = w_{\ell^*} \leq 1$.

To see that x_{i^*} satisfies (2b), first note that for $\ell > i^* + 1$, the constraint is vacuous since $x_i = 0$ for $i > i^* + 1$. For $\ell \leq i^* + 1$, the proof of the previous theorem showed that the constraint is equivalent to

$$\binom{i^*}{\ell} x_{i^*} + \binom{i^*+1}{\ell} x_{i^*+1} \leq \binom{n}{\ell} q_\ell.$$

Because $\binom{i^*}{\ell} \leq \binom{i^*+1}{\ell}$, and $w_\ell \leq w_{\ell^*}$, we have that

$$\begin{aligned}\binom{i^*}{\ell} x_{i^*} + \binom{i^*+1}{\ell} x_{i^*+1} &= \binom{i^*}{\ell} w_{\ell^*} + \binom{i^*+1}{\ell} (1 - w_{\ell^*}) \\ &\leq \binom{i^*}{\ell} w_\ell + \binom{i^*+1}{\ell} (1 - w_\ell) = \binom{n}{\ell} q_\ell.\end{aligned}$$

Finally, the objective value of x is

$$i^* x_{i^*} + (i^*+1) x_{i^*+1} = i^* + 1 - w_{\ell^*} = i^* - \frac{\binom{i^*}{\ell^*}}{\binom{i^*}{\ell^*-1}} - \frac{\binom{n}{\ell^*} q_{\ell^*}}{\binom{i^*}{\ell^*-1}},$$

where we use the fact that $\binom{i^*+1}{\ell^*} = \binom{i^*}{\ell^*} + \binom{i^*}{\ell^*-1}$. This value matches the dual bound given by Theorem 3 with $i = i^*$ and $\ell = \ell^*$. □

4 Application to Random Graph Models

4.1 Erdős-Rényi Graphs

Let $G = G_{n,p}$ be the Erdős-Rényi random graph; the independence number of this graph is well understood [2]. When p is constant, this independence number concentrates almost entirely on two numbers κ^* and $\kappa^* - 1$, where

$$\kappa^* = \min\left\{\kappa : \binom{n}{\kappa}(1-p)^{\binom{\kappa}{2}} \leq \log(n)\right\}.$$

The number κ^* is asymptotically equal to $2\log_{(1-p)^{-1}}(n) + O(\log\log(n))$ when p is constant. Similar asymptotics apply when p is at least $n^{-\delta}$ for $\delta > 0$ [8].

We consider the related quantity

$$k^* = \min\left\{k : \binom{n}{k}(1-p)^{\binom{k}{2}} \leq 1\right\};$$

k^* is also asymptotically equal to $2\log_{(1-p)^{-1}}(n) + o(\log(n))$ when p is constant.

Theorem 5 *If $\mathcal{G}$ is the probability distribution of $G_{n,p}$, then*

$$\psi(\mathcal{G}) \leq k^* + \binom{n}{k^*}(1-p)^{\binom{k^*}{2}} - 1.$$

Proof Follows from the simpler bound in Theorem 3 by setting $i = k^* - 1$. □

Theorem 5 gives the asymptotically correct answer for large n when p is a constant. Moreover, the theorem's bound is a deterministic bound on $\mathbb{E}[\alpha(G_{n,p})]$ which holds for any n and p. In contrast, when p is constant the expected value of $\vartheta(G)$ is $\Omega(\sqrt{n})$ [9], and even stronger relaxations based on $\vartheta(G)$ remain far from tight in expectation [5].

4.2 Uniformly Random Spanning Trees

Let G be a uniformly random spanning tree with vertex set $[n]$. Such random spanning trees are of interest because they can be sampled efficiently [3], and they have connections to the matrix tree theorem. The expected independence number of G was computed exactly in [15]; in particular, they showed that

$$\limsup_{n\to\infty} \frac{\mathbb{E}[\alpha(G)]}{n} = \varrho,$$

where $\varrho \approx 0.5671$ is the unique solution to the equation $xe^x = 1$. The LP (2) leads to the following bound, where we use $H(x) := -x\log(x) - (1-x)\log(1-x)$ to denote the binary entropy function.

Theorem 6 *Let $\mathcal{G}$ be the probability distribution of the uniformly random spanning tree.*

$$\limsup_{n\to\infty} \frac{\psi(\mathcal{G})}{n} \leq \varrho',$$

where $\varrho' \approx 0.6399$ is the unique solution to the equation $0 = H(x) + x\log(1-x)$ in the interval $(0,1)$.

Before proving this theorem, we establish a lemma.

Lemma 1 *The probability that the set $[\ell]$ is independent in G is precisely*

$$q_\ell = \left(1 - \frac{\ell}{n}\right)^{\ell-1}.$$

Proof Because we are choosing a uniformly random spanning tree,

$$q_\ell = \frac{\text{Number of spanning trees of } [n] \text{ so that } [\ell] \text{ is independent.}}{\text{Number of spanning trees of } [n]}$$

Kirchoff's theorem implies that the number of spanning trees of $[n]$ is n^{n-2}. We then count the number of spanning trees where $[\ell]$ is independent, using the matrix tree theorem. Spanning trees where $[\ell]$ is independent are the same as spanning trees of the graph with the edge set $\{\{i,j\} : i \notin [\ell] \text{ or } j \notin [\ell]\}$. By the matrix-tree theorem, the number of such spanning trees is precisely $\det(\tilde{L})$, where $\tilde{L}$ denotes the graph Laplacian of this graph with a row and corresponding column removed.

In this case, we may take $\tilde{L}$ to be

$$\tilde{L} = \begin{pmatrix} (n-\ell)I_{\ell-1} & -1_{\ell-1,n-\ell} \\ -1_{n-\ell,\ell-1} & nI_{n-\ell} - 1_{n-\ell,n-\ell} \end{pmatrix},$$

where $1_{a,b}$ denotes the all 1's matrix with dimensions $a \times b$, and I_ℓ denotes the $\ell \times \ell$ identity matrix.

By the Schur complement theorem for determinants,

$$\det(\tilde{L}) = \det((n-\ell)I_{\ell-1}) \det\left(nI_{n-\ell} - 1_{n-\ell,n-\ell} - \frac{\ell-1}{n-\ell} 1_{n-\ell,n-\ell}\right).$$

We obtain $\det((n-\ell)I_{\ell-1}) = (n-\ell)^{\ell-1}$, and

$$\det\left(nI_{n-\ell} - 1_{h-\ell,n-\ell} - \left(\frac{\ell-1}{n-\ell} - 1\right) 1_{n-\ell,n-\ell}\right) = n^{n-\ell-1}.$$

The last equality follows because the eigenvalues of $nI_{n-\ell} - \frac{n-1}{n-\ell} 1_{n-\ell,n-\ell}$ are n with multiplicity $n-\ell-1$ and 1 with multiplicity 1.

Thus, overall, we have that the number of spanning trees where $[\ell]$ is independent is $(n-\ell)^{\ell-1} n^{n-\ell-1}$. This implies the lemma. □

Proof (of Theorem 6) We appeal to Theorem 3 with $\ell = cn$, for some constant $0 < c < 1$ to be determined later. We have

$$\psi(\mathcal{G}) \leq cn - 1 + \binom{n}{cn} \left(1 - \frac{cn}{n}\right)^{cn-1}.$$

For large enough n, Stirling's approximation implies that

$$\binom{n}{cn} \leq C 2^{H(c)n},$$

for some absolute constant C and where $H(c)$ is the binary entropy function.

Thus, for large enough n,

$$\psi(\mathcal{G}) \leq cn - 1 + C 2^{(H(c) + \log(1-c)c)n}.$$

Choosing c so that $H(c) + \log(1-c)c = 0$, for n large enough, $\psi(\mathcal{G}) \leq cn - 1 + C$. □

4.3 Geometric Random Graphs

Let X be a metric space and let π be a probability distribution on X. The *geometric random graph* associated to π with radius r is obtained by taking n independent samples from π as the vertices and connecting any pair of samples which are at distance at most r. Such graphs have been studied in a variety of different settings, e.g. [4]. We consider the case in which X is the d-dimensional sphere with the spherical metric, and π is the uniform distribution on X. We cannot compute the probabilities q_ℓ exactly, but we have the following bound.

Lemma 2 *For the geometric random graph defined above, $q_\ell \leq e^{-V(r/2)\binom{\ell}{2}}$, where $V(r)$ is the probability of a point lying in a given ball of radius r.*

Proof We proceed by induction on ℓ. In the case $\ell = 1$, we have that $q_\ell = 1$ and the bound is trivial. For $\ell > 1$, we let $x_1, \ldots, x_\ell$ denote the points sampled for the vertices of this graph. Using Bayes' theorem,

$$\Pr\Big(\min_{i,j\leq \ell} d(x_i, x_j) > r\Big) = \Pr\Big(\min_{i,j\leq \ell-1} d(x_i, x_j) > r\Big) \cdot \Pr\Big(\min_{j\leq \ell-1} d(x_\ell, x_j) > r \Big| \min_{i,j\leq \ell-1} d(x_i, x_j) > r\Big).$$

By induction,

$$\Pr\Big(\min_{i,j\leq \ell-1} d(x_i, x_j) > r\Big) \leq e^{-V(r/2)\binom{\ell-1}{2}}.$$

We also note that

$$\Pr\Big(\min_{j\leq \ell-1} d(x_\ell, x_j) > r \Big| \min_{i,j\leq \ell-1} d(x_i, x_j) > r\Big) \leq 1-(\ell-1)V(r/2) \leq e^{-V(r/2)(\ell-1)}.$$

For this, simply note that if $\min_{i,j\leq \ell-1} d(x_i, x_j) > r$, the balls of radius $r/2$ centered at each of the x_i, for $i \leq \ell - 1$, are disjoint. If $(\ell-1)V(r/2)) > 1$, this is a contradiction, and if $(\ell-1)V(r/2)) < 1$, the probability that x_ℓ is not contained in any of the previous balls of radius $r/2$ is at most $1 - (\ell - 1)V(r/2)$, as desired. □

Theorem 7 *Let $\ell^* = \min\{\ell : \binom{n}{\ell} e^{-V(r/2)\binom{\ell}{2}} < 1\}$*

$$\psi(\mathcal{G}) \leq \ell^* - 1 + \binom{n}{\ell^*} e^{-V(r/2)\binom{\ell^*}{2}}.$$

Proof Follows directly from Theorem 3. □

5 Numerical Experiments

We begin our numerical experiments with the case $G = G_{n,p}$ when n is fixed and p varies. Figure 1 shows results for $n = 20$, comparing the first two levels of the hierarchy to the final level of the hierarchy, as well as a simulated estimate of the expected independence number $\bar{\alpha}(G)$. We obtain this estimate by sampling 100 instances and computing the deterministic independence number in each. Observe that the second-level bound is significantly better than the first-level bound, and the LP bound is very close to the empirical mean.

Our second example is non-symmetric. Given $p \in \mathbb{R}^{\binom{[n]}{2}}$, consider a random graph $G = G_p$ where the edge i, j appears independently from the others with probability p_{ij}. We will let $\mathcal{G}_p$ be the probability distribution of G_p. We let $n = 10$ and consider 100 such random models G_p, where each instance of $p \in \mathbb{R}^{\binom{[n]}{2}}$ is generated by taking independent, uniformly random numbers in the interval $[0, 1]$. For each p, we compare the bounds $\phi_t(\mathcal{G}_p)$ in the first three levels of the SDP hierarchy, as well as the empirical mean of the independence number $\bar{\alpha}(G_p)$ (calculated using 100 samples). Figure 2 summarizes the results obtained for different choices of p. The figure shows a histogram of the ratios $\phi_t(G_p)/\bar{\alpha}(G_p)$.

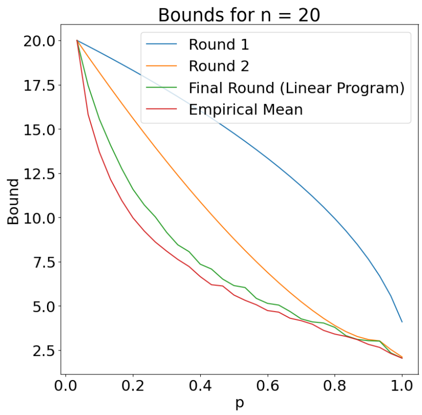

Fig. 1: The plot shows our bounds $\phi_t(\mathcal{G})$ for $t = 1, 2, n$ as well as an empirical estimate $\bar{\alpha}(G)$ of the expected independence number for the Erdős-Rényi graph $G = G_{n,p}$ with $n = 20$, for different values of p.

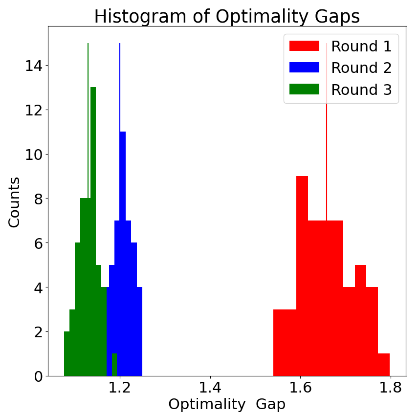

Fig. 2: A histogram of the ratios $\phi_t(\mathcal{G}_p)/\bar{\alpha}(G_p)$ for different non-symmetric random graphs G_p, where $\bar{\alpha}$ denotes an empirical estimate of the expected independence number. We display the geometric mean of the relative gaps for each round as a vertical line.

The histogram illustrates again a significant gap between the first and the second level of the relaxation. The gap between the second and third level is noticeable, but not as large. The computational cost of the SDP increases rapidly with t, so it is challenging to compute higher levels of the hierarchy. Since the random graphs G_p are non-symmetric, we cannot compute the final level directly.

Acknowledgements

The authors' work was partially supported by the Office of Naval Research, N00014-23-1-2631.

References

1. Alon, N., Spencer, J.H.: The probabilistic method. John Wiley & Sons (2016)
2. Bollobás, B.: Random graphs. Springer (1998)
3. Broder, A.Z.: Generating random spanning trees. In: FOCS. vol. 89, pp. 442–447 (1989)
4. Díaz, J., Petit, J., Serna, M.: Random geometric problems on [0, 1] 2. In: International Workshop on Randomization and Approximation Techniques in Computer Science. pp. 294–306. Springer (1998)
5. Feige, U., Krauthgamer, R.: The probable value of the Lovász-Schrijver relaxations for maximum independent set. SIAM Journal on Computing **32** (2003)
6. Grötschel, M., Lovász, L., Schrijver, A.: Polynomial algorithms for perfect graphs. In: North-Holland mathematics studies, vol. 88, pp. 325–356. Elsevier (1984)
7. Grötschel, M., Lovász, L., Schrijver, A.: Geometric algorithms and combinatorial optimization, vol. 2. Springer Science & Business Media (2012)
8. Janson, S., Luczak, T., Rucinski, A.: Random graphs. John Wiley & Sons (2011)
9. Juhász, F.: The asymptotic behaviour of Lovász' theta function for random graphs. Combinatorica **2**(2), 153–155 (1982)
10. Lasserre, J.B.: Global optimization with polynomials and the problem of moments. SIAM Journal on optimization **11**(3), 796–817 (2001)
11. Laurent, M.: A comparison of the Sherali-Adams, Lovász-Schrijver, and Lasserre relaxations for 0–1 programming. Mathematics of Operations Research **28**(3), 470–496 (2003)
12. Laurent, M., Vargas, L.F.: Exactness of Parrilo's conic approximations for copositive matrices and associated low order bounds for the stability number of a graph. Mathematics of Operations Research **48**(2), 1017–1043 (2023)
13. Lovász, L.: On the Shannon capacity of a graph. IEEE Transactions on Information theory **25**(1), 1–7 (1979)
14. Lovász, L., Schrijver, A.: Cones of matrices and set-functions and 0–1 optimization. SIAM journal on optimization **1**(2), 166–190 (1991)
15. Meir, A., Moon, J.W.: The expected node-independence number of random trees. In: Indagationes Mathematicae (Proceedings). vol. 76, pp. 335–341. Elsevier (1973)
16. Muir, C., Toriello, A.: Dynamic node packing. Mathematical Programming pp. 1–32 (2022)
17. Parrilo, P.A.: Structured semidefinite programs and semialgebraic geometry methods in robustness and optimization. California Institute of Technology (2000)
18. Rothvoß, T.: The Lasserre hierarchy in approximation algorithms. Lecture Notes for the MAPSP pp. 1–25 (2013), `https://sites.math.washington.edu/~rothvoss/lecturenotes/lasserresurvey.pdf`
19. Slot, L., Laurent, M.: Sum-of-squares hierarchies for binary polynomial optimization. Mathematical Programming **197**(2), 621–660 (2023)