

Finding cliques using few probes

Uriel Feige ^{*} David Gamarnik [†] Joe Neeman [‡]

Miklós Z. Rácz [§] Prasad Tetali [¶]

September 20, 2018

Abstract

Consider algorithms with unbounded computation time that probe the entries of the adjacency matrix of an n vertex graph, and need to output a clique. We show that if the input graph is drawn at random from $G_{n, \frac{1}{2}}$ (and hence is likely to have a clique of size roughly $2 \log n$), then for every $\delta < 2$ and constant ℓ , there is an $\alpha < 2$ (that may depend on δ and ℓ) such that no algorithm that makes n^δ probes in ℓ rounds is likely (over the choice of the random graph) to output a clique of size larger than $\alpha \log n$.

1 Introduction

Consider an algorithm (with unlimited computation time) that may make up to $q = n^\delta$ adaptive probes, to be defined later, $1 \leq \delta < 2$, to the adjacency matrix of an input graph drawn randomly from $G_{n, \frac{1}{2}}$, and needs to return a clique. What is the largest value of α (as a function of q , or of δ) such that the size of the output clique is at least $\alpha \log n$ with probability at least $\frac{1}{2}$ (over the choice of the random graph)? (All logarithms are in base 2.)

^{*}Weizmann Institute of Science; uriel.feige@weizmann.ac.il. Research supported in part by the Israel Science Foundation (grant No. 1388/16).

[†]MIT; gamarnik@mit.edu. Research supported in part by ONR grant N00014-17-1-2790

[‡]UT Austin; joeneeman@gmail.com. Research supported in part by the Alfred P. Sloan Foundation.

[§]Princeton University; mraz@princeton.edu. Research supported in part by NSF grant DMS 1811724.

[¶]Georgia Tech; tetali@math.gatech.edu. Research supported in part by NSF grants DMS-1407657 and DMS-1811935.

Observe that necessarily $\alpha \leq 2$, because with high probability the size of the largest clique in the input graph is roughly $2 \log n$ (see Remark 3).

Remark 1 *The set of q probes involves at most $2q$ vertices. If $2q < n$, then the algorithm may w.l.o.g. remove from the input graph $n - 2q$ vertices prior to making any probes (since the graph is random, it does not matter which vertices are removed), and consequently the assumption that $\delta \geq 1$ is without loss of generality. For $\delta = 2$ the algorithm can probe all entries in the adjacency matrix, and determine the largest clique. Hence we may also assume that $\delta < 2$.*

As a motivation for the types of algorithms considered in this paper, consider the following algorithm for finding a clique in the random graph. Run the greedy algorithm (iteratively choosing an arbitrary vertex and removing all its non-neighbors) until $2^{c\sqrt{\log n}}$ vertices remain, for some choice of constant c . Then exhaustively search among these for a clique of size $2c\sqrt{\log n}$. This gives a clique of size $\log n + c\sqrt{\log n}$ in time roughly n^{2c^2} . It is easy to see that this algorithm inspects $O(n)$ out of $O(n^2)$ total edges.

In the probe model, we allow the algorithms to run in superpolynomial time, but restrict the number of edges the algorithm is allowed to inspect. Specifically, we consider algorithms evolving dynamically over a certain fixed T number of steps. In the first step $t = 1$, the algorithm is allowed to choose any pair $(i_1, j_1), 1 \leq i_1 < j_1 \leq n$, and asks the status of this pair, namely whether it is an edge or not. Depending on the outcome, the algorithm selects a second pair $(i_2, j_2), 1 \leq i_2 < j_2 \leq n$, and asks the status of the pair (edge or not); it then selects (i_3, j_3) , and so on. The algorithm runs till $t = T$, may use unbounded computational time, and needs to produce a clique as large as possible.

If $T = O(n)$, then such an algorithm can produce a clique of size $\frac{3}{2} \log n$ as follows. Run the greedy algorithm until \sqrt{q} vertices remain and thereafter switch to exhaustive search. This gives a clique of size $\log \frac{n}{\sqrt{q}} + 2 \log \sqrt{q} = \log n + \frac{1}{2} \log q$. When $q = \Theta(n)$ the algorithm finds a clique of size $\frac{3}{2} \log n$ while inspecting $T = O(n)$ edges. In general, when $q = \Theta(n^{\delta/2})$ and $\delta \geq 1$, the algorithm finds a clique of size $(1 + \delta/2) \log n$ while inspecting $T = O(n^\delta)$ edges.

It is an interesting question whether there is any probing type algorithm that finds a clique larger than $(1 + \delta/2) \log n$ with $T = O(n^\delta)$. Not being able to make progress on this question, we consider the following modification of this question: what happens if the adaptiveness of the algorithm is limited, in the sense that the algorithm is run in a limited number of

stages? Specifically, we consider algorithms which can probe the entries of the adjacency matrix in a constant number of rounds, where all probes of a round are done in parallel. Precisely, fix a constant ℓ . In step $i = 1$ the algorithm selects a subset E_1 of the set of pairs $(i, j), 1 \leq i < j \leq n$, and the status of all edges in E_1 is revealed. Based on the outcome a set E_2 of pairs is selected by the algorithm and the status of edges in E_2 is revealed. This is continued until at most ℓ sets E_1, \dots, E_ℓ are created. The algorithm is limited by having

$$\sum_{1 \leq i \leq \ell} |E_i| \leq T.$$

When $T = O(n^\delta)$, this is effectively a bound $|E_i| = O(n^\delta)$, since $\ell = O(1)$.

Define $\alpha_\star = \alpha_\star(\delta, \ell)$ as the supremum over α such that there exists an algorithm that probes at most n^δ entries of the adjacency matrix of a graph drawn at random from $G_{n, \frac{1}{2}}$, with the probes being done in ℓ rounds, and outputs a clique of size at least $\alpha \log n$ with probability at least $1/2$. In this paper we prove both lower and upper bounds on $\alpha_\star(\delta, \ell)$. Since the size of the largest clique in a random graph is approximately $2 \log n$ (see Remark 3), we immediately have that $\alpha_\star \leq 2$. Our main result is an upper bound that is strictly better than this.

Theorem 2 *For every $\delta < 2$ and constant ℓ we have that $\alpha_\star(\delta, \ell) < 2$.*

We first present algorithms that find cliques in a few rounds of probes ($\ell \in \{1, 2, 3\}$) in Section 2, providing lower bounds on $\alpha_\star(\delta, \ell)$. We then prove Theorem 2 in Section 3. In fact, we prove an explicit upper bound as a function of δ and ℓ : $\alpha_\star(\delta, \ell) \leq 2 - \delta(1 - \delta/2)^\ell$, see Theorem 9.

A natural direction for future work is to improve the bounds in this paper. In particular, can one prove an upper bound (smaller than 2) that holds regardless of the number of rounds ℓ ? Finally, it would be desirable to find the actual value of $\alpha_\star(\delta, \ell)$.

Remark 3 *The size of the largest clique in a random graph is very well understood. Define $\omega_n = 2 \log n - 2 \log \log n + 2 \log e - 1$. Matula [7] showed that for any $\varepsilon > 0$, the clique number $\omega(G)$ of a random graph G drawn from $G_{n, \frac{1}{2}}$ satisfies $\lfloor \omega_n - \varepsilon \rfloor \leq \omega(G) \leq \lfloor \omega_n + \varepsilon \rfloor$ with probability tending to 1 as $n \rightarrow \infty$; see also [2]. To simplify exposition, in the following we will often neglect lower order terms and just state that the size of the largest clique is approximately $2 \log n$ with high probability (w.h.p.). In all such cases the arguments can be made precise by applying Matula's result; we leave the details to the reader.*

1.1 Related work

The problem of finding structure in a random graph by adaptively querying the existence of edges between pairs of vertices was recently introduced by Ferber, Krivelevich, Sudakov, and Vieira [4, 5]. In particular, they studied finding a Hamilton cycle [4] and finding long paths [5] in the adaptive query model. Conlon, Fox, Grinshpun, and He [3] also study this problem, which they term the Subgraph Query Problem, focusing on finding a copy of a target graph H , and in particular studying the case when H is a small (constant size) clique. All of these works focus on sparse random graphs.

In a similar vein, we study finding large cliques in an adaptive query model. The main difference between our work and those mentioned above—apart from the underlying random graph being dense or sparse—is that we study a query model where adaptiveness is limited to constantly many rounds of probes. It would be interesting to understand the effect of limited adaptiveness on finding other types of structures, such as a Hamilton cycle, long paths, or a particular target graph.

2 Algorithms

In this section we present algorithms that find cliques with one, two, and three rounds of probes. Interestingly, the algorithm for $\ell = 3$ rounds of probes described in Section 2.3 matches (to leading order) the performance of the greedy algorithm presented in Section 1, which corresponds to the *adaptive* probe model with no restriction on ℓ ($\ell = \infty$).

2.1 One round algorithm

Lemma 4 *For every $1 \leq \delta < 2$ we have that $\alpha_\star(\delta, 1) = \delta$.*

Proof. Let $q = n^\delta$. A simple one round algorithm probes all entries induced by $\sqrt{2q}$ vertices. This finds a clique of size approximately $2 \log \sqrt{2q} \simeq \log q$ (w.h.p.).

This is optimal up to lower order terms, due to the fact that in a complete graph, the pattern of q probes that maximizes the number of complete subgraphs of any given size is a clique of size $\sqrt{2q}$. (See Section 3 for details on the upper bound.) ■

2.2 Two rounds

Lemma 5 *For every $1 \leq \delta \leq \frac{6}{5}$ we have that $\alpha_*(\delta, 2) \geq \frac{4\delta}{3}$. For every $\frac{6}{5} \leq \delta < 2$ we have that $\alpha_*(\delta, 2) \geq 1 + \frac{\delta}{2}$.*

Proof. We first present the algorithm for the case when $q = \Theta(n)$.

Round 1: Probe all edges induced by a set $S \subset V$ of size $n^{1/6}$, and all edges connecting S to $T \subset V$ of size $n^{5/6}$ (where S and T are disjoint).

Round 2: Let $S' \subset S$ be a clique of largest size within S , and let $T' \subset T$ be the set of vertices within T that are neighbors of all vertices of S' . Note that S' has size approximately $\frac{1}{3} \log n$ and T' has size approximately \sqrt{n} (w.h.p.). Now probe all edges in T' , finding a clique of size approximately $\log n$ (w.h.p.). Together with S' , this gives a clique of size approximately $\frac{4}{3} \log n$ (w.h.p.).

For $q = \Theta(n^\delta)$ the same algorithm applies, with different set sizes. For $1 \leq \delta \leq \frac{6}{5}$, choose S to have size $n^{\delta/6}$ and T to have size $n^\delta/|S| = n^{5\delta/6}$. Then S' has size approximately $\frac{\delta}{3} \log n$, while T' has size approximately $n^{\delta/2}$ (w.h.p.). The largest clique in T' , together with S' , gives a clique of size approximately $\frac{4\delta}{3} \log n$ (w.h.p.). For $\frac{6}{5} \leq \delta < 2$, choose S to have size $n^{\frac{1}{2} - \frac{\delta}{4}}$ and T to have size $n \leq n^\delta/|S|$. Then S' has size approximately $(1 - \frac{\delta}{2}) \log n$, while T' has size approximately $n^{\delta/2}$ (w.h.p.). The largest clique in T' , together with S' , gives a clique of size approximately $(1 + \frac{\delta}{2}) \log n$ (w.h.p.).

■

2.3 Three rounds

Lemma 6 *For every $1 \leq \delta < 2$ we have that $\alpha_*(\delta, 3) \geq 1 + \delta/2$.*

Proof. We first present the algorithm for the case when $q = \Theta(n)$.

Round 1: Probe all edges induced by a set $S \subset V$ of size $n^{1/4}$.

Round 2: Let $S' \subset S$ be a clique of largest size within S , and note that S' has size approximately $\frac{1}{2} \log n - 2 \log \log n$ (up to an additive constant, w.h.p.). Let T be a set of $\frac{n}{\log n}$ vertices, with S and T disjoint. Probe all edges between S' and T .

Round 3: Let $T' \subset T$ be a set of \sqrt{n} vertices that are neighbors of all vertices of S' . Such a set T' exists with high probability. Probe all edges in T' , finding a clique of size approximately $\log n$ (w.h.p.). Together with S' , this gives a clique of size approximately $\frac{3}{2} \log n$ (w.h.p.).

For $q = \Theta(n^\delta)$ the same algorithm applies, with different set sizes. Choose S to have size $n^{(1-\delta/2)/2}$, in which case S' has size $(1 - \delta/2) \log n -$

$2 \log \log n$ (up to an additive constant, w.h.p.). The set T' will now have size $n^{\delta/2}$, and altogether we obtain a clique of size approximately $(1 + \delta/2) \log n$ (w.h.p.). \blacksquare

3 Upper bounds

In proving upper bounds on the size of the clique that can be found, we use the following definition.

Definition 7 *Given a graph $G = (V, E)$, a set $S \subset V$, and a parameter $0 \leq \beta \leq 1$, we say that S is β -covered if the subgraph induced by S contains at least $\beta \binom{|S|}{2}$ edges. Given positive integers n , m , and k (with $k < n$), and a parameter $0 \leq \beta \leq 1$, we let $N_{n,m,k,\beta}$ denote the maximum number of sets of size k that can be β -covered in an n vertex graph with m edges.*

The following theorem gives an upper bound on $N_{n,m,k,\beta}$.

Theorem 8 *Let $G = (V, E)$ be a graph with n vertices and m edges. Then the number $N_{n,m,k,\beta}$ of sets $S \subset V$ of size k that are β -covered satisfies:*

- $N_{n,m,k,\beta} \leq m^{(1-\sqrt{1-\beta})k+1} n^{(2\sqrt{1-\beta}-1)k+2}$ when $\beta \leq \frac{16}{25}$, and
- $N_{n,m,k,\beta} \leq m^{\sqrt{\beta}k/2+1} n^{(1-\sqrt{\beta})k+2}$ when $\beta \geq \frac{16}{25}$.

Moreover, these upper bounds are tight up to lower order multiplicative terms (when k is much smaller than m and n).

We defer the proof of Theorem 8 to Section 3.1. Using Theorem 8 we are now ready to prove Theorem 2. In fact, we state and prove a quantitative bound on $\alpha_\star(\delta, \ell)$, which implies Theorem 2. For $0 \leq \beta \leq 1$ and $1 \leq \delta < 2$ define

$$f(\beta, \delta) := \begin{cases} (2 - \delta)\sqrt{1 - \beta} + \delta - 1, & \text{if } \beta \in [0, \frac{16}{25}], \\ 1 - (1 - \frac{\delta}{2})\sqrt{\beta}, & \text{if } \beta \in [\frac{16}{25}, 1]. \end{cases} \quad (1)$$

Theorem 9 *We have that*

$$\alpha_\star(\delta, \ell) \leq \min_{\beta \in \Delta_{\ell-1}} \max_{i \in \{1, \dots, \ell\}} \frac{2f\left(\sum_{j=1}^i \beta_j, \delta\right)}{\sum_{j=i}^{\ell} \beta_j}, \quad (2)$$

where f is defined in (1) and $\Delta_{\ell-1}$ denotes the $(\ell-1)$ -simplex.

In particular, there exists a choice of $\beta \in \Delta_{\ell-1}$ in the formula above such that for every $1 \leq \delta < 2$ and constant ℓ we have that

$$\alpha_*(\delta, \ell) \leq 2 - \delta \left(\frac{2 - \delta}{2} \right)^\ell. \quad (3)$$

Proof. Let A be a deterministic (w.l.o.g., because the input is randomized) algorithm that takes ℓ rounds, for some constant ℓ , and makes $q = n^\delta$ probes in total, where $1 \leq \delta < 2$. We set $k = \alpha \log n$, for some $1 < \alpha < 2$, to be determined later as a function of ℓ and δ . We will show that A fails (w.h.p.) to find cliques of size k .

Let us fix nonnegative $\beta_1, \dots, \beta_\ell$ satisfying $\sum_{i=1}^\ell \beta_i = 1$, whose values will later be optimized as a function of δ . Consider ℓ identical copies of A . On a given input, all copies of A run in an identical fashion and all tentatively produce the same output clique K (of size k). We say that round i is *significant* if the number of probes to K in rounds 1 up to $i-1$ is at most $\sum_{j=1}^{i-1} \beta_j \binom{k}{2}$ and the number of probes to K in rounds 1 up to i is at least $\sum_{j=1}^i \beta_j \binom{k}{2}$. Given K and the sequence of probes, there must be at least one significant round; this can be proven by induction for instance. Copy i of A outputs K if the first significant round is i , and outputs nothing otherwise. This view of multiple copies of A is identical (in its output) to a single copy of A . We shall show that each copy succeeds to output a clique of size k with probability less than $1/(2\ell)$, and hence by a union bound the algorithm fails with probability at least $1/2$.

Consider now a single copy of A , say A_i . A set $S \subset V$ of size k is referred to as an i -eligible set if the number of probes to S in rounds 1 up to $i-1$ is at most $\sum_{j=1}^{i-1} \beta_j \binom{k}{2}$ and the number of probes to S in rounds 1 up to i is at least $\sum_{j=1}^i \beta_j \binom{k}{2}$. By definition, A_i is only allowed to output an i -eligible set. Note that to determine whether a set S is i -eligible, it suffices to see the answers to all probes up to round $i-1$, as this determines the sets of probes also in round i (and i -eligibility does not depend on the answers to the probes in round i). Let E_i be the event that at least one of the i -eligible sets is indeed a clique. Note that algorithm A_i produces no output unless event E_i holds, and hence the probability that A_i succeeds is bounded above by the probability of E_i . For each i -eligible set, after round $i-1$ there are at least $\sum_{j=i}^\ell \beta_j \binom{k}{2}$ pairs of vertices that have not yet been probed and hence the probability that this set is a clique is at most $2^{-\sum_{j=i}^\ell \beta_j \binom{k}{2}}$. To upper bound the number of i -eligible sets, observe that at least $\sum_{j=1}^i \beta_j \binom{k}{2}$ pairs of vertices of a given i -eligible set are probed up to round i (we do not care how

these probes are distributed among rounds 1 to i). Therefore the number of i -eligible sets is at most $N_{n,q,k,\sum_{j=1}^i \beta_j}$ (see Definition 7). A union bound thus gives us that

$$\mathbb{P}(E_i) \leq N_{n,q,k,\sum_{j=1}^i \beta_j} 2^{-\sum_{j=i}^{\ell} \beta_j \binom{k}{2}}. \quad (4)$$

By Theorem 8 we have that

$$\log N_{n,n^\delta,k,\beta} \leq \alpha f(\beta, \delta) \log^2 n + 4 \log n.$$

Consequently, using the notation $s_i := \sum_{j=1}^i \beta_j$ and taking logarithms in (4), we obtain that

$$\begin{aligned} \log \mathbb{P}(E_i) &\leq \alpha f(s_i, \delta) \log^2 n + 4 \log n - \frac{1 - s_{i-1}}{2} (\alpha^2 \log^2 n - \alpha \log n) \\ &= \alpha \left\{ f(s_i, \delta) - \frac{1 - s_{i-1}}{2} \alpha \right\} \log^2 n + \left(\frac{1 - s_{i-1}}{2} \alpha + 4 \right) \log n. \end{aligned}$$

We thus see that if $\alpha > 2f(s_i, \delta)/(1 - s_{i-1})$ then $\mathbb{P}(E_i) \rightarrow 0$ as $n \rightarrow \infty$. If this holds for every $i \in \{1, \dots, \ell\}$, this proves (2).

For $\ell = 1$, the expression (2) gives $\alpha_\star(\delta, 1) \leq \delta$, which is tight (see Lemma 4). In the following we assume that $\ell \geq 2$. For constant $\ell \geq 2$ and fixed $1 \leq \delta < 2$, the expression in (2) gives an optimization problem in $\beta \in \Delta_{\ell-1}$ to solve to obtain an explicit upper bound on $\alpha_\star(\delta, \ell)$. This is pursued in more detail in Section 3.2; in particular, the optimum is found when $\delta = 1$. Here we simply choose a particular $\beta' \in \Delta_{\ell-1}$ that implies (3) (and hence also proves Theorem 2).

Specifically, let

$$\beta'_i := r^{i-1} \varepsilon$$

for $i = 1, \dots, \ell$, where

$$r := \frac{2}{2 - \delta}$$

and, since $\sum_{i=1}^{\ell} \beta'_i = 1$, we have that $\varepsilon = (r - 1)/(r^\ell - 1)$. Define also, as above, $s'_i := \sum_{j=1}^i \beta'_j$. By (2), in order to show (3), it suffices to show that

$$\frac{2f(s'_i, \delta)}{1 - s'_{i-1}} \leq 2 - \delta r^{-\ell} \quad (5)$$

for every $i \in \{1, \dots, \ell\}$.

First, for $i = \ell$ we have that $s'_\ell = 1$ and hence $f(s'_\ell, \delta) = \delta/2$. We also have that $1 - s'_{\ell-1} = (r^\ell - r^{\ell-1}) / (r^\ell - 1)$. Hence

$$\frac{2f(s'_\ell, \delta)}{1 - s'_{\ell-1}} = \delta \frac{r^\ell - 1}{r^\ell - r^{\ell-1}} = 2 - \frac{\delta}{r^\ell - r^{\ell-1}},$$

where the second equality follows from the definition of r . Now (5) follows by dropping the $r^{\ell-1}$ term in the right hand side of the display above.

We now turn to $i \leq \ell - 1$. Note that $s'_i \leq s'_{\ell-1} = (r^{\ell-1} - 1) / (r^\ell - 1) < 1/r = (2 - \delta)/2 \leq 1/2$ and hence by the definition of f (see (1)) we have that

$$\frac{2f(s'_i, \delta)}{1 - s'_{i-1}} = \frac{2\{(2 - \delta)\sqrt{1 - s'_i} + \delta - 1\}}{1 - s'_{i-1}}.$$

Using the bound $\sqrt{1 - s'_i} \leq 1 - s'_i/2$ we obtain that

$$\frac{2f(s'_i, \delta)}{1 - s'_{i-1}} \leq 2 - \frac{(2 - \delta)s'_i - 2s'_{i-1}}{1 - s'_{i-1}} = 2 - \frac{\delta}{r^\ell - r^{i-1}}.$$

Now (5) follows by dropping the r^{i-1} term in the display above. ■

3.1 Bounding $N_{n,m,k,\beta}$: an extremal problem

In this subsection we prove Theorem 8, which gives essentially the tight upper bound on $N_{n,m,k,\beta}$ (see Definition 7), as we will show.

Given integers k and $1 \leq t \leq \binom{k}{2}$, let $M(k, t)$ denote the minimum, over all k -vertex graphs H with t edges, of the size (number of edges) of the maximum matching in H . We use the notation $M(k, \beta)$ when t is expressed as $\beta \binom{k}{2}$.

Lemma 10 *Using the notation as above, $N_{n,m,k,\beta} \leq \binom{m}{M(k,\beta)} \binom{n}{k-2M(k,\beta)}$.*

Proof. For a set S of size k to be β -covered, its induced subgraph must have a matching of size $M(k, \beta)$. We encode S by its matching edges followed by the remaining vertices in S . Given that G has m edges, there are at most $\binom{m}{M(k,\beta)}$ ways of encoding the matching edges, and then at most $\binom{n}{k-2M(k,\beta)}$ ways of choosing the remaining vertices. ■

Lemma 11 *Define*

$$\mu(k, \beta) := \min \left\{ \frac{\sqrt{\beta}}{2}, \left(1 - \sqrt{1 - \beta}\right) \right\} k. \quad (6)$$

We then have that

$$\lfloor \mu(k, \beta) \rfloor \leq M(k, \beta) \leq \lfloor \mu(k, \beta) \rfloor + 1.$$

Proof. Let H be an arbitrary graph on k vertices. It is well known (see, e.g., [6]) that if the maximum matching in H has size M , then H has a Gallai-Edmonds (GE) decomposition satisfying $c - s = k - 2M$, where c is the number of odd components C_1, \dots, C_c in the GE decomposition, s is the size of the separator set S , r is the size of the set R of remaining vertices, and edges exiting an odd component can only be connected to S . Every maximum matching matches all of R to itself, matches each vertex of S to a different odd component, and leaves exactly one unmatched vertex in every odd component that has no vertex matched to S . The value of a GE-decomposition is $c - s$, and an optimal GE-decomposition is one that maximizes $c - s$. If $c > 0$ in a GE-decomposition, then necessarily $c > s$.

Let H be a graph with k vertices and t edges for which the size of the maximum matching is $M(k, t)$. We make a *structural claim* that (at least) one of the following holds:

1. The optimal GE-decomposition for H is a clique plus a set of isolated vertices. If the clique size is even it serves as R in the GE-decomposition, whereas if its size is odd it serves as an odd component. In any case, every isolated vertex serves as an odd component, and S is empty.
2. The optimal GE-decomposition for H is a (complete) split graph, namely, a clique that serves as S , odd components that are singleton vertices, and edges between each odd component and all of S . (Note that for the split graph to be a GE-decomposition it must be that the independent set is larger than the clique, because $c - s$ needs to be positive.)
3. $M(k, t) = M(k, t + 1)$.

To prove the claim, we need to show that if the optimal GE-decomposition for H is neither a clique nor a split graph, then we can find a graph with $t + 1$ edges that has a GE-decomposition with the same value as that for H .

W.l.o.g. we may assume that $R \cup S$ is a clique, as otherwise we can add an edge to $R \cup S$ without decreasing $c - s$. Thereafter, if there are no odd components the GE-decomposition is a clique and we are done. Hence we may assume that there are odd components and $c > s$. As with the argument that $R \cup S$ is a clique, if we cannot add edges to H while preserving $c - s$, it must be that each odd component is a clique, and every vertex in every odd component is connected to every vertex in S (if S exists). Moreover, R is necessarily empty, as otherwise we can merge it with an odd component (the component remains odd because R is necessarily even), and completing that odd component to a clique we gain edges without changing $c - s$.

There remain several cases to consider:

1. There is exactly one odd component, C_1 . Then S is empty, and R was already previously assumed to be empty. Hence H is just the odd clique C_1 , and so is its GE-decomposition.
2. There are (at least) two odd components that are not singletons, say C_1 and C_2 . Then move all but one vertex from C_1 to C_2 (and remove the edges created between C_1 and C_2 , and make C_2 into a clique). The number of edges increases without decreasing $c - s$.
3. S is empty. From the previous two cases we may assume that there are at least two odd components, and at most one of them (say C_1) is a clique. As R is empty as well, then the GE-decomposition is a clique (C_1).
4. S is nonempty (and R is empty, as argued above), there are at least $s + 1$ odd components, and exactly one odd component (say C_1) is not a singleton (if C_1 is a singleton then we have a split graph, as desired).
 - (a) $|C_1| \geq c$. Merge S and s singletons into C_1 , add edges to the new C'_1 to make it a clique, and disconnect the remaining singletons from S . The GE-decomposition becomes a clique C'_1 , and the value $c - s$ does not change. The number of edges gained is $s|C_1|$, and lost is $(c - 1 - s)s$, so we strictly gained edges, as desired.
 - (b) $c > |C_1|$. Increase both c and s by one by making one vertex v of C_1 a new singleton component C'_1 , and moving one vertex u from C_1 to S . After updating the edges, we lose $|C_1| - 2$ edges of v , but gain $c - 1$ new edges to u , so the number of edges increases.

Having proved our structural claim, it remains to show the quantitative bounds. For the lower bound, it suffices to check the size of the maximum

matching in a clique and in a split graph, each with $\beta \binom{k}{2}$ edges, since the case that $M(k, t) = M(k, t + 1)$ is handled by considering $t + 1$, which only increases the computed bounds. For the upper bound it suffices to check the size of the maximum matching in graphs that are nearly a clique and nearly a split graph.

In a clique of size K the size of the maximum matching is $\lfloor \frac{K}{2} \rfloor$. Since the number of edges in the clique is $\binom{K}{2} = \beta \binom{k}{2}$, we must have $K \geq \sqrt{\beta}k$, which implies that the size of the maximum matching is at least $\lfloor \frac{\sqrt{\beta}}{2}k \rfloor$.

For the matching upper bound, let K be such that $\binom{K}{2} < \beta \binom{k}{2} \leq \binom{K+1}{2}$. The first inequality implies that $K < \sqrt{\beta}k + 1$. Now define a graph that is a clique of size K with an additional vertex that is connected to $\beta \binom{k}{2} - \binom{K}{2}$ vertices of the clique, and $k - (K + 1)$ singleton vertices. The size of the maximum matching in this graph is $\lfloor \frac{K+1}{2} \rfloor$ and by the inequality above we have that $\lfloor \frac{K+1}{2} \rfloor \leq \lfloor \frac{\sqrt{\beta}}{2}k \rfloor + 1$.

In a (complete) split graph, let c denote the size of the independent set. The number of non-edges is $\binom{c}{2} = (1 - \beta) \binom{k}{2}$ and hence $c \leq \lceil \sqrt{1 - \beta}k \rceil$. The size of the maximum matching is $k - c$, since in the maximum matching every vertex of the clique is matched to a vertex in the independent set, and we have that $k - c \geq \lfloor (1 - \sqrt{1 - \beta})k \rfloor$.

For the matching upper bound, let c be such that $\binom{c}{2} \leq (1 - \beta) \binom{k}{2} < \binom{c+1}{2}$. The second inequality implies that $c \geq \lceil \sqrt{1 - \beta}k \rceil - 1$. We define a graph on k vertices as follows. Start with a split graph on k vertices with an independent set of size c . Now take a vertex v from the clique of the split graph and remove $(1 - \beta) \binom{k}{2} - \binom{c}{2}$ edges connecting v to vertices of the independent set. The resulting graph has $\beta \binom{k}{2}$ edges and its maximum matching has size $k - c$. By the inequality above we have that $k - c \leq \lfloor (1 - \sqrt{1 - \beta})k \rfloor + 1$. ■

Proof.[of Theorem 8] By Lemma 10 and a simple bound on binomial coefficients we obtain that

$$N_{n,m,k,\beta} \leq \binom{m}{M(k,\beta)} \binom{n}{k - 2M(k,\beta)} \leq m^{M(k,\beta)} n^{k - 2M(k,\beta)}.$$

Using Lemma 11 we thus obtain that

$$N_{n,m,k,\beta} \leq m^{\mu(k,\beta)+1} n^{k - 2\mu(k,\beta)+2},$$

where recall the definition of $\mu(k, \beta)$ from (6). Theorem 8 now follows by observing that $\mu(k, \beta) = (1 - \sqrt{1 - \beta})k$ when $\beta \leq \frac{16}{25}$ and $\mu(k, \beta) = \frac{\sqrt{\beta}}{2}k$ when $\beta \geq \frac{16}{25}$.

This bound is tight (up to factors of $k^{\Theta(k)}$) as can be seen by the following examples. First, let $\beta \leq \frac{16}{25}$. Suppose first that $m < kn$. Consider m edges that form a split graph with a clique K of size $(1 - \sqrt{1 - \beta})k$ joined to an independent set I of size roughly $\frac{m}{k}$. Any choice of $\sqrt{1 - \beta}k$ vertices from I completes together with K a split graph with $\beta \binom{k}{2}$ edges. If $m > kn$, then make K of size m/n and I of size $n - |K|$. Any choice of $\sqrt{1 - \beta}k$ vertices from I and $(1 - \sqrt{1 - \beta})k$ vertices from K forms a split graph with $\beta \binom{k}{2}$ edges.

Now let $\beta \geq \frac{16}{25}$. Consider m edges that form a clique of size $\sqrt{2m}$. Any choice of $\sqrt{\beta}k$ vertices from the clique and $(1 - \sqrt{\beta})k$ vertices from the rest of the graph gives a subgraph with $\beta \binom{k}{2}$ edges. ■

3.2 Explicit upper bounds for $\delta = 1$

The expression (2) in Theorem 9 gives an optimization problem to compute an upper bound on $\alpha_\star(\delta, \ell)$. As mentioned in the proof of Theorem 9, for $\ell = 1$ this gives $\alpha_\star(\delta, 1) \leq \delta$, which is tight. Here we investigate the optimization problem of (2) for other values of ℓ . In particular, we solve this optimization problem for every ℓ when $\delta = 1$.

For $\ell = 2$, the expression (2) gives

$$\alpha_\star(\delta, 2) \leq \min_{\beta \in [0, 1]} \max \left\{ 2f(\beta, \delta), \frac{\delta}{1 - \beta} \right\}.$$

For the two expressions in the display above to be less than 2, we must have $\delta/(1 - \beta) < 2$, or equivalently, $\beta < 1 - \delta/2 \leq 1/2$. Hence recalling the definition of f (see (1)) we obtain that

$$\alpha_\star(\delta, 2) \leq \min_{\beta \in [0, \frac{1}{2}]} \max \left\{ 2 \left((2 - \delta) \sqrt{1 - \beta} + \delta - 1 \right), \frac{\delta}{1 - \beta} \right\}.$$

Observe that the first expression in the display above is decreasing in β , while the second expression is increasing in β . Therefore the unique minimizer $\beta' = \beta'(\delta)$ satisfies

$$2 \left((2 - \delta) \sqrt{1 - \beta'} + \delta - 1 \right) = \frac{\delta}{1 - \beta'}$$

and we have that

$$\alpha_\star(\delta, 2) \leq \frac{\delta}{1 - \beta'}.$$

For $\delta = 1$ we obtain that $\beta'(1) = 1 - 2^{-2/3}$ and thus $\alpha_\star(1, 2) \leq 2^{2/3} < 1.588$.

Many of the observations made for the case of $\ell = 2$ above also apply for $\ell \geq 3$. We may write the expression (2) as follows:

$$\alpha_\star(\delta, \ell) \leq \min_{\beta \in \Delta_{\ell-1}} \max \left\{ \max_{i \in \{1, \dots, \ell-1\}} \frac{2f(s_i, \delta)}{1 - s_{i-1}}, \frac{\delta}{\beta_\ell} \right\},$$

where $s_i = \sum_{j=1}^i \beta_j$ as before. If $\beta' = \beta'(\delta)$ denotes the optimizer then we must have $\beta'_\ell > \delta/2 \geq 1/2$. This implies that $s'_i \leq s'_{\ell-1} \leq 1/2$ for every $i \leq \ell - 1$, where $s'_i = \sum_{j=1}^i \beta'_j$. Hence recalling the definition of f again we obtain that

$$\alpha_\star(\delta, \ell) \leq \min_{\beta \in \Delta_{\ell-1}, \beta_\ell \geq \frac{\delta}{2}} \max \left\{ \max_{i \in \{1, \dots, \ell-1\}} \frac{2((2-\delta)\sqrt{1-s_i} + \delta - 1)}{1 - s_{i-1}}, \frac{\delta}{\beta_\ell} \right\}.$$

This simplifies when $\delta = 1$:

$$\alpha_\star(1, \ell) \leq \min_{\beta \in \Delta_{\ell-1}, \beta_\ell \geq \frac{1}{2}} \max \left\{ \max_{i \in \{1, \dots, \ell-1\}} \frac{2\sqrt{1-s_i}}{1 - s_{i-1}}, \frac{1}{\beta_\ell} \right\}.$$

The optimizer $\beta' = \beta'(1)$ is such that all ℓ expressions on the right hand side of the display above are equal. That is, we must have

$$\frac{2\sqrt{1-s'_i}}{1 - s'_{i-1}} = \frac{1}{\beta'_\ell}$$

for every $i \in \{1, \dots, \ell - 1\}$. This set of equations can be solved and we obtain the following optimizer:

$$\begin{aligned} \beta'_1 &= 1 - 2^{-\frac{2}{2^\ell-1}}, \\ \beta'_i &= 2^{-\frac{2^i-2}{2^\ell-1}} - 2^{-\frac{2^{i+1}-2}{2^\ell-1}} \quad \text{for } i \in \{2, \dots, \ell - 1\}, \\ \beta'_\ell &= 2^{-\frac{2^\ell-2}{2^\ell-1}}. \end{aligned}$$

The optimum is therefore $1/\beta'_\ell = 2^{\frac{2^\ell-2}{2^\ell-1}}$. In conclusion, we have proved the following corollary of Theorem 9.

Corollary 12 *We have that*

$$\alpha_\star(1, \ell) \leq 2^{1 - \frac{1}{2^\ell-1}}.$$

Numerically, the bound above gives the following for small ℓ : $\alpha_*(1, 2) \leq 2^{2/3} < 1.588$, $\alpha_*(1, 3) \leq 2^{6/7} < 1.812$, $\alpha_*(1, 4) \leq 2^{14/15} < 1.910$, and $\alpha_*(1, 5) \leq 2^{30/31} < 1.956$. These upper bounds should be compared to the lower bounds $\alpha_*(1, 2) \geq 4/3$ and $\alpha_*(1, \ell) \geq \alpha_*(1, 3) \geq 3/2$ for $\ell \geq 3$ (see Section 2).

Acknowledgements

The problem considered here was proposed by David Gamarnik at the American Institute of Mathematics workshop “*Phase transitions in randomized computational problems*” in June 2017 [1]. It arose from a discussion with Madhu Sudan, whose contribution to the inception of the problem is gratefully acknowledged. We thank AIM and the organizers, Amir Dembo, Jian Ding, and Nike Sun, for putting together the workshop. We also thank Jane Gao for initial discussions.

References

- [1] AimPL. Phase transitions in randomized computational problems. Available at <http://aimpl.org/phaserandom>, 2017.
- [2] B. Bollobás and P. Erdős. Cliques in random graphs. *Mathematical Proceedings of the Cambridge Philosophical Society*, 80(3):419–427, 1976.
- [3] D. Conlon, J. Fox, A. Grinshpun, and X. He. Online Ramsey Numbers and the Subgraph Query Problem. Preprint available at <https://arxiv.org/abs/1806.09726>, 2018.
- [4] A. Ferber, M. Krivelevich, B. Sudakov, and P. Vieira. Finding Hamilton cycles in random graphs with few queries. *Random Structures & Algorithms*, 49(4):635–668, 2016.
- [5] A. Ferber, M. Krivelevich, B. Sudakov, and P. Vieira. Finding paths in sparse random graphs requires many queries. *Random Structures & Algorithms*, 50(1):71–85, 2017.
- [6] L. Lovász and M. D. Plummer. *Matching Theory*. American Mathematical Society, 2009.
- [7] D. W. Matula. The Employee Party Problem. In *Notices of the American Mathematical Society*, volume 19, pages A–382, 1972.