

The Strongish Planted Clique Hypothesis and Its Consequences

Pasin Manurangsi

Google Research, Mountain View, CA, USA
pasin@google.com

Aviad Rubinfeld

Stanford University, CA, USA
aviad@cs.stanford.edu

Tselil Schramm

Stanford University, CA, USA
tselil@stanford.edu

Abstract

We formulate a new hardness assumption, the *Strongish Planted Clique Hypothesis (SPCH)*, which postulates that any algorithm for planted clique must run in time $n^{\Omega(\log n)}$ (so that the state-of-the-art running time of $n^{O(\log n)}$ is optimal up to a constant in the exponent).

We provide two sets of applications of the new hypothesis. First, we show that SPCH implies (nearly) tight inapproximability results for the following well-studied problems in terms of the parameter k : Densest k -Subgraph, Smallest k -Edge Subgraph, Densest k -Subhypergraph, Steiner k -Forest, and Directed Steiner Network with k terminal pairs. For example, we show, under SPCH, that no polynomial time algorithm achieves $o(k)$ -approximation for Densest k -Subgraph. This inapproximability ratio improves upon the previous best $k^{o(1)}$ factor from (Chalermsook et al., FOCS 2017). Furthermore, our lower bounds hold even against fixed-parameter tractable algorithms with parameter k .

Our second application focuses on the complexity of graph pattern detection. For both induced and non-induced graph pattern detection, we prove hardness results under SPCH, improving the running time lower bounds obtained by (Dalirrooyfard et al., STOC 2019) under the Exponential Time Hypothesis.

2012 ACM Subject Classification Theory of computation → Problems, reductions and completeness; Theory of computation → Fixed parameter tractability

Keywords and phrases Planted Clique, Densest k -Subgraph, Hardness of Approximation

Digital Object Identifier 10.4230/LIPIcs.ITCS.2021.10

Related Version A full version of the paper is available at <https://arxiv.org/abs/2011.05555>.

Funding *Aviad Rubinfeld*: A.R. is supported in part by a Packard Fellowship.

Tselil Schramm: T.S. is supported by a Simons Institute Microsoft Research Fellowship.

Acknowledgements Pasin would like to thank Michal Pilipczuk and Daniel Lokshtanov for posing the approximability of Densest k -Subhypergraph as an open problem at Dagstuhl Seminar on New Horizons in Parameterized Complexity, and for helpful discussions.

1 Introduction

The last couple of decades have seen dramatic advances in our understanding of parameterized, fine-grained, and average-case complexity. To a large extent, this progress has been enabled by bolder computational hardness assumptions, beyond the classical $P \neq NP$. Two notable assumptions in these fields are the Exponential Time Hypothesis and the Planted Clique Hypothesis. In this paper we propose a new hypothesis, the *Strongish Planted Clique*



© Pasin Manurangsi, Aviad Rubinfeld, and Tselil Schramm;
licensed under Creative Commons License CC-BY

12th Innovations in Theoretical Computer Science Conference (ITCS 2021).

Editor: James R. Lee; Article No. 10; pp. 10:1–10:21

Leibniz International Proceedings in Informatics



LIPICs Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany

Hypothesis, which strengthens the Planted Clique Hypothesis in the style of the Exponential Time Hypothesis. We show that this hypothesis has interesting implications in the parameterized complexity of both approximation problems and graph pattern detection.

Exponential Time Hypothesis. The Exponential Time Hypothesis (ETH) [69] is a pessimistic version of $P \neq NP$ which postulates that solving 3-SAT on n variables requires time $2^{\Omega(n)}$. In other words, the running times of the state-of-the-art (and the brute-force) algorithms for 3-SAT are optimal up to a constant factor in the exponent. ETH has important applications in parameterized complexity (e.g. [30, 82, 80, 46, 48]) and hardness of approximation (e.g. [26, 25, 92, 84, 87, 52]). In the past several years, further progress was achieved by assuming stronger variants of ETH such as the Strong ETH (SETH) in fine-grained complexity [70, 95], Gap-ETH in parameterized complexity [31, 23, 83], and ETH for PPAD in Algorithmic Game Theory [9, 91].

Planted Clique Hypothesis. In the Planted κ -Clique Problem, the goal is to distinguish (with high probability) between graphs sampled from one of the following distributions: Uniformly at random¹; and uniformly at random, with an added κ -clique. While statistically it is easy to distinguish the two distributions for κ as little as $2.1 \log(n)$, the *Planted Clique Hypothesis* (PCH) postulates that no polynomial time algorithm can solve this problem, even for κ as large as $o(\sqrt{n})$. The history of this problem goes back to Karp [73] and Jerrum [71], and in the past decade it has been popular as a hardness assumption for both worst-case [8, 68, 3, 10, 20] and average-case [17, 66, 18, 63, 29, 27] problems. A simple $n^{\Theta(\log(n))}$ -time algorithm for the planted- κ -clique problem non-deterministically guesses $\ell = \Theta(\log(n))$ vertices from the clique, and then checks whether all of their common neighbors form a clique. There are several other algorithms that also solve this problem in time $n^{\Theta(\log(n))}$ [60, 59, 86, 13], but no faster algorithm is known for $\kappa = O(n^{0.49})$.

Strongish Planted Clique Hypothesis

In analogy to the Exponential Time Hypothesis for 3-SAT, we propose the following hypothesis, which postulates that the state-of-the-art algorithms for the Planted Clique Problem are optimal up to a constant factor in the exponent. A *Strong* Planted Clique Hypothesis, in analogy with SETH, would specify a precise constant in the exponent – our hypothesis is merely Strong-ish. We let $\mathcal{G}(n, p)$ denote the Erdős-Rényi distribution with parameter p , and $\mathcal{G}(n, p, \kappa)$ denote the Erdős-Rényi distribution with a planted κ -clique.

► **Hypothesis 1** (Strongish Planted Clique Hypothesis (SPCH)). *There exists a constant $\delta \in (0, \frac{1}{2})$ such that no $n^{o(\log n)}$ -time algorithm \mathcal{A} satisfies both of the following:*

- (Completeness) $\Pr_{G \sim \mathcal{G}(n, \frac{1}{2}, \lceil n^\delta \rceil)}[\mathcal{A}(G) = 1] \geq 2/3$.
- (Soundness) $\Pr_{G \sim \mathcal{G}(n, \frac{1}{2})}[\mathcal{A}(G) = 1] \leq 1/3$.

In addition to the lack of algorithmic progress toward refuting this hypothesis, we note that $n^{\Theta(\log(n))}$ is in fact provably optimal for the Sum-of-Squares hierarchy [11], which captures the state-of-the-art algorithmic techniques for a number of average-case problems. It is also known to be tight for statistical algorithms [62, 28].

¹ I.e. from the *Erdős-Rényi* (ER) distribution over n -vertex graphs where each edge appears independently with probability $1/2$.

1.1 Our Contributions: Hardness results from Strongish Planted Clique Hypothesis

Our main technical contributions are in exploring the implications of our new SPCH in parameterized complexity. We prove two types of hardness results: hardness of approximation, and hardness of (exact) graph pattern detection. Due to the rich literature for each problem we consider, we will only mention the most relevant results here and defer a more comprehensive discussion to Section 1.4.

1.1.1 Hardness of approximation from SPCH

At the heart of our work is the study of the *Densest k -Subgraph* problem, in which we are given an undirected graph $G = (V, E)$ and a positive integer k . The goal is to output a subset $S \subseteq V$ of k vertices that induces as many edges as possible. There is a trivial $O(k)$ -approximation for the problem: return an arbitrary set of $\lfloor k/2 \rfloor$ edges. We show that assuming SPCH, this algorithm is optimal:

► **Theorem 2.** *Assuming the Strongish Planted Clique Hypothesis (Hypothesis 1), there is no $f(k) \cdot \text{poly}(n)$ -time algorithm that can approximate Densest k -Subgraph on n -vertex graphs to within a factor $o(k)$ for any function f . Furthermore, this holds even in the perfect completeness case where the input graph is promised to contain a k -clique.*

Theorem 2 improves upon the inapproximability ratio of $k^{o(1)}$ shown in [31] under Gap-ETH.

The approximability of Densest k -Subgraph is known to be intimately related to that of numerous other problems. As such, our tight hardness of approximation for Densest k -Subgraph immediately implies several tight approximability results as corollaries, which we list below.

- *Smallest k -Edge Subgraph*: given an undirected graph $G = (V, E)$ and a positive integer k , find a smallest subset $S \subseteq V$ that induces at least k edges. For this problem, the trivial solution that chooses k edges arbitrarily is an $O(\sqrt{k})$ -approximation since even the optimum requires at least \sqrt{k} vertices. We show that this is tight (Corollary 11): no fixed-parameter tractable (FPT) (in k) algorithm can achieve $o(\sqrt{k})$ approximation ratio.
- *Steiner k -Forest* (aka *k -Forest*): given an edge-weighted undirected graph $G = (V, E)$, a set $\{(s_1, t_1), \dots, (s_\ell, t_\ell)\}$ of demand pairs and a positive integer k , the goal is to find a (not necessarily induced) subgraph of G with smallest total edge weight that connects at least k demand pairs. In Corollary 12, we show that no FPT (in k) algorithm can achieve $o(\sqrt{k})$ approximation ratio. This matches the $O(\sqrt{k})$ -approximation algorithm by Gupta et al. [65].
- *Directed Steiner Network* (aka *Directed Steiner Forest*): given an edge-weighted directed graph $G = (V, E)$ and a set $\{(s_1, t_1), \dots, (s_k, t_k)\}$ of k demand pairs, the goal is to find a (not necessarily induced) subgraph of G with smallest total edge weight in which t_i is reachable from s_i for all $i = 1, \dots, k$. We prove that no FPT (in k) algorithm achieves $o(\sqrt{k})$ -approximation (Corollary 14). This nearly matches the best known polynomial algorithms by Chekuri et al. [35] and Feldman et al. [61], both of which achieve a $O(k^{1/2+\epsilon})$ -approximation for any constant $\epsilon > 0$. Our bound improves upon a $k^{1/4-o(1)}$ ratio from [52] under Gap-ETH.
- *Densest k -Subhypergraph*: given a hypergraph $G = (V, E)$ and a positive integer k , output a k -size subset $S \subseteq V$ that maximizes the number of hyperedges fully contained in S . The trivial algorithm that outputs any hyperedge (of size at most k) obtains a 2^k -approximation. We prove a matching lower bound (Theorem 15): no $2^{o(k)}$ -approximation FPT (in k) algorithm exists. This resolves an open question posed by Cygan et al. [45], assuming SPCH.

1.1.2 Hardness of graph pattern detection from SPCH

Graph Pattern Detection, also known as *Subgraph Isomorphism* and closely related to *Motif Discovery*, is a fundamental problem in graph algorithms: Given a host graph G and a pattern graph H , decide whether G contains a subgraph S isomorphic to H . There are two main variants of this problem, where S is either required to be *induced* or *not-necessarily-induced*. For both variants, there is a brute-force $n^{O(k)}$ -time algorithm. We prove matching SPCH-hardness for both variants. In addition to beating the ETH-based state-of-the-art for these results, we highlight that our reductions to graph pattern detection problems are extremely simple (in contrast to the prior work).

For induced subgraph detection, we prove the following:

► **Theorem 3.** *Assuming the Strongish Planted Clique Hypothesis (Hypothesis 1), for every k -node pattern H , there is no algorithm that solves the induced pattern detection problem on n -vertex graphs in time $f(k) \cdot n^{o(k)}$ for any function f .*

Our $n^{\Omega(k)}$ lower bound for all patterns can be compared to recent work of [48], who proved: (i) $n^{\Omega(\log(k))}$ lower bound for every pattern assuming ETH; (ii) $n^{\Omega(\sqrt{k})}$ lower bound for every pattern assuming ETH and the Hadwiger conjecture; and (iii) $n^{\Omega(k/\log(k))}$ lower bound for most patterns.²

For not-necessarily-induced subgraph detection, it is no longer true that every pattern is hard (e.g. it is trivial to find a not-necessarily-induced subgraph isomorphic to an independent set). But we prove (Corollary 10) that for most patterns³ detection requires $n^{\Omega(k)}$ time assuming SPCH. For comparison, [48] proved that under ETH, not-necessarily-induced subgraph detection requires $n^{\Omega(\omega(H))}$ time, where $\omega(H)$ denotes the clique number of the pattern H . (Note that $\omega(H) = \Theta(\log k)$ for most patterns.)

k -biclique detection. For the special case where the pattern is a k -biclique, our aforementioned $n^{\Omega(k)}$ hardness for non-induced subgraph detection rules out even constant factor approximations (Corollary 9). This improves over $n^{\Omega(\sqrt{k})}$ lower bounds under ETH for the exact case [80] or Gap-ETH for approximation [31].

Densest k -Subgraph. We obtain our pattern detection result by first showing a $n^{\Omega(k)}$ running time bound for $O(1)$ -approximating Densest k -Subgraph (Theorem 8). This improves upon the previous lower bound who give a running time lower bound of $n^{\Omega(\log k)}$ assuming Gap-ETH [31]. The aforementioned lower bounds for pattern detection follow almost trivially from our running time lower bound for Densest k -Subgraph; see Sections 3.2 and 3.3 for more detail.

1.2 Techniques

The starting point for all of our reductions is a *randomized graph product*: starting with an instance G of the planted clique problem on n vertices and any integers $\ell \leq n$ and N , we produce a graph G' by taking its vertices to be N randomly sampled subsets S_1, \dots, S_N of ℓ vertices each, and we add an edge on S_i, S_j if and only if their union induces a clique in G .

² In the sense of a pattern sampled randomly from $\mathcal{G}(k, 1/2)$.

³ in particular any pattern with a constant fraction of the $\binom{k}{2}$ possible edges.

The randomized graph product [16] (and its derandomized variant [4]) has a long history in proving hardness of approximating Maximum Clique. While not stated explicitly, it was also used to prove parameterized inapproximability of Densest k -Subgraph in [31]. As we will explain in more detail below, the main difference between our proof and previous works lie in the soundness, where we appeal to the fact that $G \sim \mathcal{G}(n, 1/2)$ to achieve tighter bounds.

Since we would like to prove hardness of approximating Densest k -Subgraph in the perfect completeness case, our goal is to show (for appropriately chosen values of N, ℓ) that

- (Completeness) if G contains a large clique, then with high probability so does G' , and,
- (Soundness) if G is random, then G' does not have small dense subgraphs (with high probability).

For the completeness, if the starting graph G has a κ -clique, then the set of S_i that fall entirely within the κ -clique will form a clique in G' (the expected size is $(\frac{\kappa}{n})^\ell \cdot N$). This part of the proof is exactly the same as that in the aforementioned previous works.

To prove soundness, we calculate the probability that a specific γk -dense k -subgraph appears in G' , then take a union bound over all $\leq \binom{N}{k} \cdot 2^{\binom{k}{2}}$ possible such subgraphs. Our simple argument hinges on showing that a k -subgraph with γk^2 edges in G' induces (with high probability) at least $\Omega(\gamma k^2 \ell^2)$ edges in G , and since any such set of edges appears in G with probability at most $2^{-\Omega(\gamma k^2 \ell^2)}$, by choosing ℓ sufficiently large we can beat this union bound. To argue that small subgraphs with m edges in G' induce small subgraphs with $\Omega(m \ell^2)$ edges in G , we use that the randomly chosen S_i (for an appropriate choice of $N \ll n^\ell$ and k_* sufficiently small) form a *dispenser*: the union of *any* $t \leq k_*$ of the S_i contains $\Omega(t \ell)$ vertices of G with high probability.⁴ This implies that for $k \leq k_*$, each k -subgraph of G' corresponds to a union of k *pairwise* mostly-non-overlapping subsets of ℓ vertices. Now, since each edge between mostly non-overlapping sets in G' corresponds to a $\Omega(\ell)$ -clique in G , this in turn can be used to show that any k -subgraph of G' with density γk corresponds to a subgraph of G with $\Omega(\gamma k^2 \ell^2)$ edges. In this way we rule out the existence of γk -dense k -subgraphs in G' (with high probability).

By carefully choosing the parameters N, ℓ, γ , to control the completeness, soundness, and reduction size, we get a fine-grained reduction from Planted κ -Clique to Densest k -Subgraph. Our results for other problems are obtained via direct reductions from the Densest k -Subgraph problem.

We end by stressing that our new soundness proof gives a strong quantitative improvement over prior results, which is what enables us to achieve $k^{\Omega(1)}$ inapproximability. Specifically, the previous soundness proof from [31] – in turn adapted from [84] – relies on showing that the graph is t -biclique-free for some $t \in \mathbb{N}$; they then apply the so-called Kovari-Sos-Turan theorem [77] to deduce that any k -subgraph contains at most $O(k^{2-1/t})$ edges. Notice that this gap is only $O(k^{1/t})$, and t cannot be a constant as otherwise the completeness and soundness case can be distinguished in time $n^{O(t)} = \text{poly}(n)$. As a result, their technique cannot yield an $k^{\Omega(1)}$ -factor inapproximability for Densest k -Subgraph. Similarly, the techniques from [31] cannot give a running time lower bound of the form $n^{\omega(\log k)}$ for $O(1)$ -approximation of Densest k -Subgraph. The reason is that, to get a constant gap bounded from one, they need to select $t = O(\log k)$. Once again, this is in contrast to our technique which yields a tight running time lower bound of $n^{\Omega(k)}$ in this setting, although our proof requires a different starting hardness result (from SPCH).

⁴ The fact that the randomized graph product yields a dispenser has been used in previous works as well, see e.g. [16, 98, 31].

1.3 Discussion and Open Questions

In this work, we have proposed the Strongish Planted Clique Hypothesis, and used it to prove several tight hardness of approximation or running time lower bound results. One direction for future investigation is to use SPCH to derive other interesting lower bounds.

Another intriguing question directly related to our hardness of approximation results is whether we can strengthen the inapproximability factors from $k^{\Omega(1)}$ to $n^{\Omega(1)}$. Whether it is hard to approximate Densest k -Subgraph to within a $n^{\Omega(1)}$ factor is a well-known open problem and is related to some other conjectures, such as the Sliding-Scaling Conjecture [14]⁵. As such, it would be interesting if one can prove this hardness under some plausible assumption. We remark that an attempt has been made in this direction, using the so-called “Log-Density Threshold” approach [21], which posits a heuristic for predicting which average-case Densest- k -Subgraph problems are hard. The approach has also been applied to other related questions [39, 40, 42]. Nonetheless, there is still little evidence that these average-case DkS problems are indeed hard; not even lower bounds against strong SDP relaxations are known, although there are some matching lower bounds against LP hierarchies [22, 41].

Finally, it is of course interesting to either refute or find more evidence supporting the Strongish Planted Clique Hypothesis. As stated earlier, the current best supporting evidence is the Sum-of-Squares lower bound from [11]. Can such a lower bound be extended to, e.g., rule out any semi-definite programs of size $n^{o(\log n)}$ (à la [79] for CSPs)?

1.4 Other Related Works

Historically, postulating hardness for average-case problems has been helpful in illuminating the landscape for hardness of approximation, beginning with Feige’s seminal random-3-SAT hypothesis [57] and its numerous consequences (e.g. [50]). See also [49, 2, 7, 12].

As discussed briefly above, the Planted Clique Hypothesis (PCH), which states that there is no polynomial-time algorithm for planted clique, has many known consequences for hardness of approximation. We draw attention in particular to the work of [3], which also obtains hardness of approximation results based on PCH. But even assuming the SPCH, their results can only rule out $n^{\text{polylog}(\kappa)}$ -time algorithms for $2^{\log(\kappa)^{2/3}}$ -approximating densest- κ -subgraph for $\kappa = n^{\Omega(1)}$. Their reduction also uses a graph product, but the set of vertices of their new graph G' contains all ℓ -size subsets of vertices of G . In contrast, we employ the *randomized* graph product, where we only randomly pick some ℓ -size subsets – this allows us to better control the instance size blowup, which turns out to be crucial for obtaining our tight inapproximability and running time results.

Below we discuss in more detail the previous works for each of the problems studied here.

Densest k -Subgraph. The problem is well-studied in the approximation algorithms literature (e.g. [58, 76, 21]). The best known polynomial time algorithm [21] achieves an approximation ratio of $O(n^{1/4+\epsilon})$ for any $\epsilon > 0$. Even though the NP-hardness of Densest k -Subgraph follows immediately from that of k -Clique [74], no NP-hardness of approximation of Densest k -Subgraph even for a small factor of 1.001 is known. Nonetheless, several hardness of approximation results are known under stronger assumptions [57, 75, 90, 3, 12, 25, 84].

⁵ Specifically, from a reduction of [34], $n^{\Omega(1)}$ -factor hardness of approximation of Densest k -Subgraph also implies that of the Label Cover problem.

Among these, only [3, 12] and [84] yield super-constant inapproximability ratios. Specifically, [3] rules out $2^{O((\log n)^{2/3})}$ -approximation in polynomial time under a hypothesis similar to SPCH, [84] rules out $n^{o(1/\text{poly log log } n)}$ -approximation under ETH, and [12] rules out $n^{O(1)}$ approximations under a strong conjecture regarding the optimality of semidefinite programs for solving every random CSP.

Our hardness result holds even in the so-called *perfect completeness* case, where we are promised that the graph G contains a k -clique. In this case, a quasi-polynomial time approximation scheme exists [60]. Braverman et al. [25] showed that this is tight: there exists a constant $\epsilon > 0$ for which $(1 + \epsilon)$ -approximation of Densest k -Subgraph in the perfect completeness case requires $n^{\Omega(\log n)}$ -time assuming ETH. We remark that the hardness from [84] also applies in the perfect completeness case, but it only rules out polynomial time algorithms.

For the parameterized version of the problem, its W[1]-hardness follows immediately from that of k -Clique [54]. Chalermsook et al. [31] showed that no $k^{o(1)}$ -approximation is achievable in FPT time, unless Gap-ETH is false. This hardness also holds in the perfect completeness case, and yields a running time lower bound of $n^{\Omega(\log k)}$ for any constant factor approximation.

k -Biclique. Similar to Densest k -Subgraph, while the NP-hardness for the exact version of k -Biclique has long been known (e.g. [72]), even 1.001-factor NP-hardness of approximation has not yet been proved although inapproximability results under stronger assumptions are known [57, 75, 19, 85]. Specifically, under strengthened variants of the Unique Games Conjecture, the problem is hard to approximate to within a factor of $n^{1-\epsilon}$ for any $\epsilon > 0$ [19, 85].

On the parameterized complexity front, whether k -Biclique is FPT (in k) was a well-known open question (see e.g. [55]). This was resolved by Lin [80] who showed that the problem is W[1]-hard and further provided a running time lower bound of $n^{\Omega(\sqrt{k})}$ under ETH. As stated above, this running time lower bound was extended to rule out any constant factor approximation in [31] under Gap-ETH. Furthermore, [31] showed that no $o(k)$ -approximation exists in FPT time.

Smallest k -Edge Subgraph. Most of the hardness of approximation for Densest k -Subgraph easily translates to Smallest k -Edge Subgraph as well, with a polynomial loss in the factor of approximation. For example, the hardness from [84] implies that Smallest k -Edge Subgraph cannot be approximated to within a factor of $n^{1/\text{poly log log } n}$ in polynomial time, assuming ETH. On the other hand, Chlamtac et al. [39] devised an $O(n^{3-2\sqrt{2}+\epsilon})$ -approximation algorithm for any constant $\epsilon > 0$ for the problem; this remains the best known approximation algorithm for the problem.

Densest k -Subhypergraph. Apart from the hardness results inherited from Densest k -Subgraph, not much is known about Densest k -Subhypergraph. Specifically, the only new hardness is that of Applebaum [6], who showed that the problem is hard to approximate to within n^ϵ for some constant $\epsilon > 0$, assuming a certain cryptographic assumption; this holds even when each hyperedge has a constant size. On the other hand, the only (non-trivial) approximation algorithm is that of Chlamtac et al. [37] which achieves $O(n^{0.698})$ -approximation when the hypergraph is 3-uniform.

Steiner k -Forest. The Steiner k -Forest problem is a generalization of several well-known problems, including the Steiner Tree problem and the k -Minimum Spanning Tree problem. This problem was first explicitly defined in [67] and subsequently studied in [93, 65, 51].

In terms of the number of vertices n of the input graph, the best known approximation ratio achievable in polynomial time is $O(\sqrt{n})$ [65] (assuming that $k \leq \text{poly}(n)$); furthermore, when edge weights are uniform, a better approximation ratio of $O(n^{0.449})$ is achievable in polynomial time [51]. On the other hand, as stated earlier, in terms of k , the best known approximation ratio is $O(\sqrt{k})$ [65].

A reduction in [67] together with W[1]-hardness of k -Clique [54] implies that Steiner k -Forest is W[1]-hard with respect to the parameter k . We are not aware of any further parameterized complexity study of this problem (with respect to parameter k).

Directed Steiner Network. Several polynomial time approximation algorithms have been proposed for the Directed Steiner Network problem [33, 35, 61, 15, 38, 1]; in terms of the number of vertices n of the input graph the best known approximation ratio is $O(n^{2/3+\epsilon})$ [15] and in terms of k the best known ratio is $O(k^{1/2+\epsilon})$ for any constant $\epsilon > 0$ [35, 61]. On the hardness front, Dodis and Khanna [53] shows that the problem is quasi-NP-hard to approximate to within a factor of $2^{(\log n)^{1-\epsilon}}$ for any constant $\epsilon > 0$. Furthermore, Guo et al. [64] show that the exact version of the problem is W[1]-hard with respect to parameter k . Later, [52] rules out even $k^{1/4-o(1)}$ -approximation in FPT time, under Gap-ETH.

Graph Pattern Detection. As discussed earlier, [48] give ETH-based hardness results for graph pattern detection, both in the induced and non-induced case. The complexity for many special patterns has also been considered, e.g. k -cliques, k -bicliques (mentioned above), and k -cycles (e.g. [5, 97, 47, 48, 81]). A k -clique can be detected in time $O(n^{\lceil k/3 \rceil \omega})$ using fast matrix multiplication [89], and the k -Clique Conjecture in Fine-Grained Complexity postulates that this is essentially optimal [95]. Any other pattern over k vertices can be detected in time $O(n^{k-1})$, without using fast matrix multiplication [24]. There is also an extensive body of work on *counting* the number of occurrences of a pattern in a host graph (e.g. [88, 78, 94, 96, 44, 43]).

Preliminaries and Notation

For a natural number $n \in \mathbb{N}$, we use $[n]$ to denote the set of integers up to n , $[n] = \{1, \dots, n\}$. We will use the abbreviation “w.h.p.” to mean “with high probability.”

For an undirected graph $G = (V, E)$, we use $\deg_G(v)$ to denote the degree of a vertex $v \in V$, and $\min\text{-deg}(G)$ to denote $\min_{v \in V(G)} \deg_G(v)$. For a subset $S \subseteq V$, we use $G[S] = (V, E[S])$ to denote the induced subgraph of G on subset S .

The *density* of G , denoted by $\text{den}(G)$, is defined as $|E|/|V|$. We use $\text{den}_{\leq k}(G)$ to denote $\max_{S \subseteq V, |S| \leq k} \text{den}(G[S])$, the maximum density of subgraphs of G of at most k vertices.

2 Randomized Graph Product

In this section we formally define our reduction, and analyze its soundness and completeness in terms of the reduction parameters (in later sections we instantiate the parameters differently for each target bound). Our reduction takes as input a graph and applies the *randomized graph product* [16], described in Figure 1. We use $\text{RP}_{N,\ell}(G)$ to denote the distribution of outputs of the above reduction on input graph G .

We will show that for well-chosen N and ℓ , the randomized graph product $\text{RP}_{N,\ell}$ reduces the planted n^δ -clique problem to densest- k subgraph for $k = k(N, \ell, \delta, n)$. That is, a sample from $\text{RP}_{N,\ell} \circ \mathcal{G}(n, \frac{1}{2})$ has no dense k -subgraph with probability close to 1, and if on the other hand G is a graph with an n^δ -clique, then a sample from $\text{RP}_{N,\ell}(G)$ has a dense k -subgraphs with probability close to 1.

Input: n -vertex Graph $G = (V, E)$, positive integers N, ℓ .
Output: Graph $G' = (V', E')$.
The graph G' is constructed as follows.

1. For each $i \in [N]$, sample $S_i \subseteq V$ by independently sampling ℓ vertices uniformly from V .
2. Let $V' = \{S_1, \dots, S_N\}$.
3. For every distinct $i, j \in [N]$, include (S_i, S_j) in E' iff $S_i \cup S_j$ induces a clique in G .

■ **Figure 1** Randomized Graph Product [16].

2.1 Completeness

We first prove that applying the randomized graph product to a graph with a large clique results in a graph with a large clique.

► **Lemma 4 (Completeness).** *Suppose that $\delta \in (0, 1), N, \ell, k \in \mathbb{N}$ are such that $N \geq 10k \cdot n^{(1-\delta)\ell}$ and $k \geq 20$. If G contains a clique of size $\lceil n^\delta \rceil$, then*

$$\Pr_{G' \sim \text{RP}_{N, \ell}(G)} [G' \text{ contains a } k\text{-clique}] \geq 0.9.$$

Proof. Let $C \subseteq V$ be the $\lceil n^\delta \rceil$ -size clique in G . For each $i \in [N]$, $\Pr[S_i \subseteq C] = \left(\frac{\lceil n^\delta \rceil}{n}\right)^\ell \geq n^{-(1-\delta)\ell}$. By our lower bound on N , the expected number of indices $i \in [N]$ such that $S_i \subseteq C$ is at least $10k$, and thus a Chernoff bound implies that with probability $1 - \exp(-4k) \geq 0.9$, there exists at least k indices $i \in [N]$ such that $S_i \subseteq C$. By definition of the randomized graph product $\text{RP}_{N, \ell}$, these subsets form a clique in G' . ◀

2.2 Soundness

We now prove that if we apply the randomized graph product to a graph drawn from $\mathcal{G}(n, \frac{1}{2})$, with high probability the resulting graph has no small subgraphs which are too dense.

► **Lemma 5 (Soundness).** *Suppose that $\delta \in (0, 1), N, \ell, k \in \mathbb{N}$ are such that $N \leq 1000k \cdot n^{(1-\delta)\ell}$, $\ell \geq k \geq 20$ and $n^{0.99\delta} \geq k\ell$. If G is drawn from $\mathcal{G}(n, \frac{1}{2})$, then*

$$\Pr_{\substack{G \sim \mathcal{G}(n, \frac{1}{2}) \\ G' \sim \text{RP}_{N, \ell}(G)}} \left[\text{den}_{\leq k}(G') \leq \frac{10^7 \log n}{\ell \delta^2} \right] \geq 0.9.$$

We will use the following observation that allows us to translate a large-density graph to a subgraph with large minimum degree. This observation is folklore and appears e.g. in [32].

► **Observation 6.** *For any $H = (V_H, G_H)$, there exists $S' \subseteq V_H$ such that $\text{min-deg}(H[S']) \geq \text{den}(H)$.*

Another auxiliary lemma that is useful for us is that for any subset $M \subseteq [N]$ not too large, the size of the union $|\cup_{j \in M} S_j|$ is not too small relative to $|M| \cdot \ell$. This lemma is also standard and has been used in prior works (e.g. [16, 98, 31]).

► **Lemma 7.** *Suppose $N \leq 1000\ell n^{(1-\delta)\ell}$, $20 \leq \ell$. Then with probability at least 0.95 over a sample $G' \sim \text{RP}_{N, \ell} \circ \mathcal{G}(n, \frac{1}{2})$, $G' = (\{S_i\}_{i \in [N]}, E')$, the following event occurs: for every $M \subseteq [N]$ with $|M| \leq n^{0.99\delta}/\ell$, we have $|\cup_{i \in M} S_i| \geq 0.01\delta|M|\ell$.*

10:10 The Strongish Planted Clique Hypothesis and Its Consequences

The proofs of both Observation 6 and Lemma 7 are included in the full version. We now prove the soundness guarantee.

Proof of Lemma 5. We will assume that the event in Lemma 7 occurs and show that conditioned on this event, under the randomness of G , with probability 0.95 all k -subgraphs of G' have density at most $d := 10^7 \log n / (\ell \delta^2)$. Lemma 5 then follows immediately since $(0.95)^2 > 0.9$.

Consider any $J \subseteq [N]$ of size $k' \leq k$. For brevity, let $\mathcal{F}(J)$ denote the set of all graphs whose vertices are S_j for $j \in J$ and the minimum degree is at least d . For each $F = (\{S_j\}_{j \in J}, E_F) \in \mathcal{F}(J)$, let $\mathcal{E}^G(F)$ denote $\bigcup_{\{S_j, S_{j'}\} \in E_F} \{\{u, v\} \mid u \in S_j, v \in S_{j'} \text{ and } u \neq v\}$, or in words, the set of edges of G which have one endpoint in S_j and one endpoint in $S_{j'}$ for an edge $(S_j, S_{j'}) \in E_F$. Observe that

$$\Pr_G[G'[\{S_j\}_{j \in J}] = F] \leq \Pr_G[\mathcal{E}^G(F) \subseteq E] = 2^{-|\mathcal{E}^G(F)|}, \quad (1)$$

where the inequality follows by definition of the randomized graph product – since (S_i, S_j) is an edge if and only if $S_i \cup S_j$ is a clique in G , the event $G'[\{S_j\}_{j \in J}] = F$ contains the event $\mathcal{E}^G(F) \subseteq E$ – and the final equality follows because $G \sim \mathcal{G}(n, \frac{1}{2})$. To bound $|\mathcal{E}^G(F)|$, let $\mathcal{V}^G(F) := \bigcup_{j \in J} S_j$.

Since we have conditioned on the event in Lemma 7 occurring,⁶ we have $|\mathcal{V}^G(F)| \geq 0.01\delta k' \ell$. Now, consider any $v \in \mathcal{V}^G(F)$; from definition of $\mathcal{V}^G(F)$, $v \in S_j$ for some $j \in J$. Since $F \in \mathcal{F}(J)$, S_j must have at least d neighbors in F . Let $S_{j_1}, \dots, S_{j_{d'}}$ denote S_j 's neighbors in F , with $d' \geq d$. By applying the bound in Lemma 7, we have $|S_{j_1} \cup \dots \cup S_{j_{d'}}| \geq 0.01\delta d' \ell \geq 0.01\delta d \ell$. In other words, v has degree at least $0.01\delta d \ell - 1 \geq 0.005\delta d \ell$ in the graph $(\mathcal{V}^G(F), \mathcal{E}^G(F))$. This implies that

$$|\mathcal{E}^G(F)| \geq \frac{1}{2} (0.01\delta k' \ell) (0.005\delta d \ell) \geq 10^{-5} \delta^2 k' d \ell^2.$$

Plugging the above bound back into (1), we have

$$\Pr_G[G'[\{S_j\}_{j \in J}] = F] \leq 2^{-10^{-5} \delta^2 k' d \ell^2} \quad (2)$$

We can use the above inequality to bound the probability that $\{S_j\}_{j \in J}$ induces a subgraph with minimum degree at least d as follows:

$$\begin{aligned} \Pr[\min\text{-deg}(G'[\{S_j\}_{j \in J}]) \geq d] &= \sum_{F \in \mathcal{F}(J)} \Pr_G[G'[\{S_j\}_{j \in J}] = F] \stackrel{(2)}{\leq} 2^{(k')^2} \cdot 2^{-10^{-5} \delta^2 k' d \ell^2} \\ &\leq 2^{-10^{-6} \delta^2 k' d \ell^2}. \end{aligned}$$

where the first inequality follows because there are at most $2^{(k')^2}$ subgraphs of an k' -vertex graph, and to obtain the final inequality we have applied that $\ell \geq k \geq k'$ and $d \geq 10^7 / (\delta^2 \ell)$.

Applying Observation 6, the existence of a k -subgraph with density at least d would imply the existence of some $J \subseteq [N]$ with $|J| \leq k$ and minimum degree at least d . Taking union bound over all $J \subseteq [N]$ of size at most k and applying our above bound, we have

$$\Pr[\text{den}_{\leq k}(G)] \leq \sum_{k'=1}^k N^{k'} \cdot 2^{-10^{-6} \delta^2 k' d \ell^2} = \sum_{k'=1}^k \left(N \cdot 2^{-10^{-6} \delta^2 d \ell^2} \right)^{k'}$$

⁶ Note that $|J| = k' \leq k \leq n^{0.995} / \ell$ by our assumption and hence J satisfies the condition in Lemma 7.

$$\begin{aligned} &\leq \sum_{k'=1}^k \left(8 \cdot n \cdot 2^{-10^{-6}\delta^2 d \ell}\right)^{k' \ell} \\ &\leq \sum_{k'=1}^k 0.01^{k' \ell} \leq 0.95, \end{aligned}$$

where to obtain the second line we use that $N \leq 1000kn^{(1-\delta)\ell}$ and that $\ell \geq k \geq 20$, and in the final line we use that $d \geq 10^7 \log n / (\ell\delta)^2$. This completes our proof. \blacktriangleleft

3 Tight Running Time Lower Bounds

In this section, we prove our tight running time lower bounds for $O(1)$ -approximating Densest k -Subgraph, $O(1)$ -approximating k -Biclique and (exact) Graph Pattern Detection.

3.1 Constant Approximation for Densest k -Subgraph

We start with the $n^{\Omega(k)}$ running time lower bound for $O(1)$ -approximating Densest k -Subgraph, from which the remaining results easily follow. We remark that this running time lower bound improves upon that of $n^{\Omega(\log k)}$, which is implicit in [31].

► Theorem 8. *Assuming Hypothesis 1, for any constant $C > 0$, there is no $f(k) \cdot n^{o(k)}$ -time algorithm that can approximate Densest k -Subgraph to within a factor C for any function f . Furthermore, this holds even in the perfect completeness case where the input graph is promised to contain a k -clique.*

We will prove Theorem 8 by simply selecting an appropriate setting of parameters for the randomized graph product. Specifically, we will let $\ell = O(C \log n / k)$ and $N = n^{O(\ell)} = n^{o(\log n)}$; the generic soundness lemma (Lemma 5) then implies that the density of any k -subgraph in the soundness case is at most $O(k/C)$ which yields the desired C inapproximability factor.

Proof of Theorem 8. We will reduce the problem of distinguishing samples from $\mathcal{G}(n, \frac{1}{2})$ vs. $\mathcal{G}(n, \frac{1}{2}, \lceil n^\delta \rceil)$ to approximating Densest k -subgraph.

For C the constant specified in the statement of the theorem, choose $\ell = \lceil \frac{10^8 C \log n}{\delta^2 k} \rceil$ and $N = \lceil 100kn^{(1-\delta)\ell} \rceil$, so that $d = \frac{10^7 \log n}{\ell \delta^2} \leq \frac{k}{10C}$. Given a graph G , we sample $G' \sim \text{RP}_{N,\ell}(G)$.

Completeness: If $G \sim \mathcal{G}(n, \frac{1}{2}, \lceil n^\delta \rceil)$, then by Lemma 4 and since $N \geq 10kn^{(1-\delta)\ell}$, we have that with probability at least $1 - \exp(-4k) \geq 0.9$, $G' = \text{RP}_{N,\ell}(G)$ has a k -clique, and so $\text{den}_{\leq k}(G') \geq k - 1$.

Soundness: For any $20 \leq k \leq \ell$ and any δ bounded away from 0, we clearly satisfy the requirement of Lemma 5 that $k\ell \leq n^{0.99\delta}$ and that $N \leq 1000kn^{(1-\delta)\ell}$ for any sufficiently large n . Hence, if $G \sim \mathcal{G}(n, \frac{1}{2})$, applying the Lemma to $G' \sim \text{RP}_{N,\ell}(G)$ we have that with probability at least 0.9, $\text{den}_{\leq k}(G') \leq d \leq \frac{k}{10C}$.

Thus, any algorithm which approximates Densest k -subgraph up to a factor of C in time $f(k)N^{\epsilon k}$ can solve the planted $\lceil n^\delta \rceil$ -clique problem in time $f(k)(100kn^{(1-\delta)\ell})^{\epsilon k} = g(k)n^{\frac{(1-\delta)\delta^2 C}{10^8} \epsilon \log n}$ for $g(k) = f(k)(100k)^{\epsilon k}$. This contradicts Hypothesis 1 whenever $\lim_{n \rightarrow \infty} \epsilon = 0$. Choosing k to be a sufficiently slowly growing function of n , for any ϵ decreasing in k we have a contradiction of the Hypothesis. This concludes the proof. \blacktriangleleft

3.2 $O(1)$ -Approximation for k -Biclique

Recall that a k -biclique, denoted by $K_{k,k}$, is a complete bipartite graph where each side has exactly k vertices. In the k -Biclique problem, we are given an undirected graph G and a positive integer k . Further, we are promised that G contains a $K_{k,k}$ as a subgraph. Our goal is to output a balanced biclique in G of size as large as possible. (Note that we say that an algorithm achieves α -approximation if the output biclique has size at least k/α .)

We prove the following tight running time lower bound for $O(1)$ -approximation of k -Biclique:

► **Corollary 9.** *Assuming Hypothesis 1, for any constant $C > 0$, there is no $f(k) \cdot n^{o(k)}$ -time algorithm that can approximate k -Biclique to within a factor C for any function f . Furthermore, this holds even when we are promised that the input graph contains a $2k$ -clique.*

Proof. Suppose contrapositively that there is an $f(k) \cdot n^{o(k)}$ -time algorithm \mathbb{A} that can approximate k -Biclique to within a factor of C for some function f . We may use it to approximate Densest k -Subgraph with perfect completeness as follows. On a given instance (G, k) of Densest k -Subgraph with perfect completeness, we run our algorithm \mathbb{A} on $(G, \lfloor k/2 \rfloor)$. From the approximation guarantee of \mathbb{A} , we will find a t -Biclique for $t \geq \lfloor k/2 \rfloor / C \geq \frac{k}{4C}$. By taking this biclique together with arbitrary $(k - 2t)$ additional vertices, we find a subset of size k that induces at least $t^2 \geq \frac{k^2}{16C^2}$ edges. Hence, we have found a $16C^2$ -approximation for Densest k -Subgraph in time $f(k) \cdot n^{o(k)}$, which by Theorem 8 violates Hypothesis 1. ◀

3.3 Pattern detection

Theorem 8 also yields the following corollary for the running time of pattern detection:

► **Corollary 10.** *Assuming Hypothesis 1, for almost all k -node patterns H , the not necessarily induced pattern detection problem cannot be solved in time $f(k) \cdot n^{o(k)}$.*

In the statement of Corollary 10, by “almost all k -node patterns” we mean w.h.p. over $H \sim \mathcal{G}(k, 1/2)$.⁷

Proof. By standard concentration inequalities, most $H \sim \mathcal{G}(k, \frac{1}{2})$ have average degree $\frac{k}{2} \pm o(k)$. For such a pattern H , an algorithm that solves the not necessarily induced pattern detection problem also gives $O(1)$ -approximation for Densest k -Subgraph. Hence, the corollary immediately follows from Theorem 8. ◀

► **Theorem (Restatement of Theorem 3).** *Assuming the Strongish Planted Clique Hypothesis (Hypothesis 1), for every k -node pattern H , there is no algorithm that solves the induced pattern detection problem on n -vertex graphs in time $f(k) \cdot n^{o(k)}$ for any function f .*

Proof. We assume that H is at least $\frac{k}{4}$ -dense (that is, has average degree at least $\frac{k}{2}$). This is without loss of generality as otherwise, we may take the complement of H – for induced pattern detection the complexity is the same.

We start our reduction from an instance G of Densest k -Subgraph as in Theorem 8. We randomly color all the vertices of G in k colors, one for each vertex of H . We construct a graph G' from G by keeping edge $(u, v) \in G$ if and only if u, v have different colors, and the vertices in H corresponding to those colors share an edge. (Note that we do not add any edges.)

⁷ It may be more natural to sample uniformly from all *unlabeled* k -node patterns, but w.h.p. a graph drawn from $\mathcal{G}(k, 1/2)$ has no non-trivial automorphisms [56], so the two notions of “almost all k -node patterns” are in fact equivalent.

Completeness: If G has a k -clique, with probability at least k^{-k} it has k different colors. In this case, the same vertices in G' will form an induced copy of H .

Soundness: If G does not have a $\frac{k}{4}$ -dense k -subgraph, this remains true after removing edges.

Hence G' also does not contain any $\frac{k}{4}$ -dense k -subgraph – and in particular no copy of H . As a result, if we can solve the induced pattern detection problem in time $f(k) \cdot n^{o(k)}$, we can achieve 4-approximation for Densest k -Subgraph with probability k^{-k} in time $f(k) \cdot n^{o(k)}$. By repeating this construction $100 \cdot k^k$ times, we can achieve 4-approximation for Densest k -Subgraph with probability 0.99 in time $O(f(k) \cdot k^k) \cdot n^{o(k)}$. Together with Theorem 8, this concludes our proof. ◀

4 Tight Inapproximability Results

In this section, we prove tight inapproximability results for Densest k -Subgraph, Smallest k -Edge Subgraph, Steiner k -Forest, Directed Steiner Network, and Densest k -Subhypergraph.

4.1 Densest k -Subgraph

We start with the $o(k)$ -factor hardness of Densest k -Subgraph (Theorem 2), from which our other results follow. The proof of Theorem 2 is once again via selecting an appropriate setting of parameters for the randomized graph product. Specifically, we will select $\ell = O((\log n) \cdot g(k)/k)$ where $g(k) = o(k)$ is the assumed approximation ratio and $N = n^{O(\ell)} = n^{o(\log n)}$; the generic soundness lemma (Lemma 5) then implies that the density of any k -vertex subgraph in the soundness case is at most $k/g(k)$ as desired. The arguments are formalized below.

Proof of Theorem 2. We will reduce the problem of distinguishing samples from $\mathcal{G}(n, \frac{1}{2})$ vs. $\mathcal{G}(n, \frac{1}{2}, \lceil n^\delta \rceil)$ to approximating Densest k -subgraph on an N vertex graph.

Let $g(k) \leq k$ be any function growing with k . Choose $\ell = \lceil \frac{10^8 g(k) \log n}{\delta^2 k} \rceil$ and $N = \lceil 100kn^{(1-\delta)\ell} \rceil$, so that $d = \frac{10^7 \log n}{\ell \delta^2} \leq \frac{k}{10g(k)}$. Given a graph G , we sample $G' \sim \text{RP}_{N,\ell}(G)$.

Completeness: If $G \sim \mathcal{G}(n, \frac{1}{2}, \lceil n^\delta \rceil)$, then just as in the proof of Theorem 8, Lemma 4 implies that with probability at least 0.9, $G' = \text{RP}_{N,\ell}(G)$ has a k -clique.

Soundness: For any $20 \leq k \leq \ell$ and δ bounded away from 0, we satisfy the requirements of Lemma 5 (just as in the proof of Theorem 8). Applying the lemma, if $G \sim \mathcal{G}(n, \frac{1}{2})$, we conclude that with probability at least 0.9, $\text{den}_{\leq k}(G') \leq d \leq \frac{k}{10g(k)}$. This means that any k -vertex subgraph of G' contains at most $\frac{k^2}{10g(k)} < \frac{\binom{k}{2}}{g(k)}$ edges.

Hence, any algorithm which approximates Densest k -subgraph within $g(k)$ in time $f(k) \text{poly}(N)$ can solve the planted $\lceil n^\delta \rceil$ -clique problem in time $f(k) \cdot \text{poly}(100kn^{(1-\delta)\ell}) = h(k) \cdot \text{poly}\left(n^{\frac{(1-\delta)\delta^2}{10^8} \frac{g(k)}{k} \log n}\right)$ for $h(k) = f(k) \text{poly}(100k)$. Choosing k to be a sufficiently slowly growing function of n , for any function $g(k)$ with $\lim_{n \rightarrow \infty} \frac{g(k)}{k} = 0$ we have a contradiction of Hypothesis 1, as desired. ◀

4.2 Smallest k -Edge Subgraph

► **Corollary 11.** *Assuming Hypothesis 1, there is no $f(k) \cdot \text{poly}(n)$ -time algorithm that can approximate Smallest k -Edge Subgraph to within a factor $o(\sqrt{k})$ for any function f .*

We prove the above corollary by reducing from Densest k -Subgraph; we remark here that similar reductions between the two problems are folklore and have appeared before in literature, e.g. in [67]. Hence, we defer the proof to the full version of the paper.

4.3 Steiner k -Forest

► **Corollary 12.** *Assuming Hypothesis 1, there is no $f(k) \cdot \text{poly}(n)$ -time algorithm that can approximate Steiner k -Forest to within a factor $o(\sqrt{k})$ for any function f .*

Corollary 12 immediately follows from Corollary 11 via the following reduction:

► **Theorem 13** ([67, Theorem 6.2]). *If there is an $f(k) \cdot \text{poly}(n)$ -time $g(k)$ -approximation algorithm for Steiner k -Forest, then there is an $f(k) \cdot \text{poly}(n)$ -time $g(k)$ -approximation algorithm for Smallest k -Edge Subgraph.*

4.4 Directed Steiner Network

► **Corollary 14.** *Assuming Hypothesis 1, there is no $f(k) \cdot \text{poly}(n)$ -time algorithm that can approximate Directed Steiner Network with k terminal pairs to within a factor $o(\sqrt{k})$ for any function f .*

Corollary 14 is shown via a reduction from Smallest p -Edge Subgraph, which is similar to the reduction from the Label Cover problem by Dodis and Khanna [53] that was also used in subsequent works [36, 52]. The proof is deferred to the full version.

4.5 Densest k -Subhypergraph

► **Theorem 15.** *Assuming Hypothesis 1, there is no $f(k) \cdot \text{poly}(n)$ -time algorithm that can approximate Densest k -Subhypergraph to within a factor $2^{o(k)}$ for any function f .*

The proof of our inapproximability for Densest k -Subhypergraph (Theorem 15) is unlike the others in this section: instead of starting from the inapproximability of Densest k -Subgraph (Theorem 2), we will start from the tight running time lower bound for $O(1)$ -approximate k -Biclique (Corollary 9).

4.5.1 A Combinatorial Lemma

Before we describe our reduction, let us prove the following lemma, which bounds the number of (induced) copies of $K_{\ell,\ell}$ in a $K_{t,t}$ -free graph for $\ell < t$. This is a generalized setting of the classic Kovari-Sos-Turan (KST) theorem [77], which applies only to the case $\ell = 1$. Note that, due to the parameters of interest in our reduction, we only prove a good bound for large ℓ ; for $\ell = 1$, the bound in our lemma is trivial (and hence weaker than the KST theorem).

► **Lemma 16.** *Let κ, t, ℓ be positive integers such that $\ell < t \leq \kappa/16$. Then, for any κ -vertex $K_{t,t}$ -free graph H , the number of (not necessarily induced) copies of $K_{\ell,\ell}$ in H is at most $\left(2 \cdot e^{-\frac{\ell^2}{16t}}\right) \cdot \binom{\kappa}{\ell} \binom{\kappa-\ell}{\ell}$.*

Proof. Let V denote the vertex set of H . We will count the number of copies of $K_{\ell,t}$ in H in two ways. First, for every subset $S \in \binom{V}{\ell}$, the number of (ℓ, t) -biclques of the form (S, T) where $T \in \binom{V}{t}$ is $\binom{|N(S)|}{t}$ where $N(S)$ denote the set of common neighbors of S . Hence, in total the number of (ℓ, t) -biclques in H is $\sum_{S \in \binom{V}{\ell}} \binom{|N(S)|}{t}$. On the other hand, for every set $T \in \binom{V}{t}$, we must have $|N(T)| \leq t - 1$ since H does not contain $K_{t,t}$. As a result, the number of (ℓ, t) -biclques of the form (S, T) for a fixed T is at most $\binom{t-1}{\ell}$. That is, in total, there can be at most $\binom{\kappa}{t} \binom{t-1}{\ell}$ copies of $K_{\ell,t}$ in the graph. This implies that

$$\binom{\kappa}{t} \binom{t-1}{\ell} \geq \sum_{S \in \binom{V}{\ell}} \binom{|N(S)|}{t}. \quad (3)$$

For brevity, let λ denote $(\kappa - \ell)/t$, γ denote $\lambda^{-\ell/(2t)}$ and let x denote $\gamma \cdot \kappa + (1 - \gamma) \cdot t$ (note x is chosen so that $\frac{x-t}{\kappa-t} = \gamma$). Let us now classify $S \in \binom{A}{\ell}$ into two groups: those with $|N(S)| \geq x$ and those with $|N(S)| < x$. That is, we define

$$\mathcal{S}_{\geq x} := \left\{ S \in \binom{V}{\ell} \mid |N(S)| \geq x \right\}, \quad \text{and} \quad \mathcal{S}_{< x} := \left\{ S \in \binom{V}{\ell} \mid |N(S)| < x \right\}.$$

From (3), we have $\binom{\kappa}{t} \binom{t-1}{\ell} \geq \sum_{S \in \mathcal{S}_{\geq x}} \binom{|N(S)|}{t} \geq |\mathcal{S}_{\geq x}| \cdot \binom{\lceil x \rceil}{t}$. Rearranging, we have

$$\begin{aligned} |\mathcal{S}_{\geq x}| &\leq \frac{\binom{t-1}{\ell} \binom{\kappa}{t}}{\binom{\lceil x \rceil}{t}} \leq \binom{t-1}{\ell} \left(\frac{\kappa-t}{x-t} \right)^t = \binom{t-1}{\ell} \gamma^{-t} \leq \binom{\kappa}{\ell} \left(\frac{t-1}{\kappa} \right)^\ell \gamma^{-t} \\ &\leq \binom{\kappa}{\ell} \lambda^{-\ell} \gamma^{-t} \\ &= \binom{\kappa}{\ell} \lambda^{-\ell/2} \leq \binom{\kappa}{\ell} 2^{-\ell/2}, \end{aligned} \quad (4)$$

where the last line follows because $t, \ell \leq \frac{1}{3}\kappa$.

Let us now count the number of $K_{\ell, \ell}$ in H . For each fixed $S \in \binom{V}{\ell}$, the number of $K_{\ell, \ell}$ of the form (S, T) where $T \in \binom{V}{\ell}$ is exactly $\binom{|N(S)|}{\ell}$. Thus, the total number of $K_{\ell, \ell}$ in G is $\sum_{S \in \binom{V}{\ell}} \binom{|N(S)|}{\ell}$. This term can be further bounded by

$$\begin{aligned} \sum_{S \in \binom{V}{\ell}} \binom{|N(S)|}{\ell} &= \sum_{S \in \mathcal{S}_{\geq x}} \binom{|N(S)|}{\ell} + \sum_{S \in \mathcal{S}_{< x}} \binom{|N(S)|}{\ell} \\ &\leq |\mathcal{S}_{\geq x}| \binom{\kappa - \ell}{\ell} + |\mathcal{S}_{< x}| \binom{x}{\ell} \\ &\stackrel{(4)}{\leq} 2^{-\ell/2} \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} + \binom{\kappa}{\ell} \binom{\lceil x \rceil}{\ell} \\ &\leq 2^{-\ell/2} \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} + \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} \cdot \left(\frac{x}{\kappa - \ell} \right)^\ell \\ &\leq 2^{-\ell/2} \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} + \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} \cdot \left(1 - \frac{(1 - \gamma)(\kappa - t) - \ell}{\kappa - \ell} \right)^\ell \\ &= 2^{-\ell/2} \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} + \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} \cdot \left(1 - \frac{(1 - \gamma - \frac{\ell}{\kappa - t})(\kappa - t)}{\kappa - \ell} \right)^\ell \\ (\text{From } \ell < t \leq \kappa/2) &\leq 2^{-\ell/2} \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} + \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} \cdot (1 - 0.5(1 - \gamma - 2\ell/\kappa))^\ell \\ &\leq 2^{-\ell/2} \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} + \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} \cdot e^{-0.5\ell(1 - \gamma - 2\ell/\kappa)} \end{aligned} \quad (5)$$

Consider the term $(1 - \gamma - 2\ell/\kappa)$. We can bound it as follows:

$$\begin{aligned} (1 - \gamma - 2\ell/\kappa) &= \left(1 - \frac{1}{\lambda^{\ell/2t}} - \frac{2\ell}{\kappa} \right) \\ (\text{From } \lambda = (\kappa - \ell)/t \geq 2) &\geq \left(1 - 0.5^{\ell/2t} - \frac{2\ell}{\kappa} \right) \\ (\text{Bernoulli inequality}) &\geq \left(1 - \left(1 - \frac{\ell}{4t} \right) - \frac{2\ell}{\kappa} \right) = \left(\frac{\ell}{4t} - \frac{2\ell}{\kappa} \right) \geq \frac{\ell}{8t}. \end{aligned}$$

Plugging this back into (5), we have

$$\sum_{S \in \binom{A}{\ell}} \binom{|N(S)|}{\ell} \leq \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} \cdot \left(2^{-\ell/2} + e^{-\frac{\ell^2}{16t}}\right) \stackrel{(\text{From } \ell < t)}{\leq} \binom{\kappa}{\ell} \binom{\kappa - \ell}{\ell} \cdot \left(2 \cdot e^{-\frac{\ell^2}{16t}}\right),$$

which concludes our proof. \blacktriangleleft

4.5.2 Proof of Theorem 15

Proof of Theorem 15. Suppose contrapositively that there is an $f(\rho) \cdot \text{poly}(n)$ -time $2^{\rho/g(\rho)}$ -approximation algorithm \mathbb{A} for Densest ρ -Subhypergraph where $g = \omega(1)$. We will construct an algorithm \mathbb{B} that achieves $O(1)$ -approximation for k -Biclique with the promise that a $2k$ -Clique exists in time $h(k) \cdot n^{o(k)}$ for some function h . Theorem 15 then follows from Corollary 9.

We define \mathbb{B} on input $(G = (V, E), k)$ as follows:

- Let $\rho = 2k$ and $\ell = \lceil \rho/g(\rho)^{0.1} \rceil$.
- Construct a hypergraph G' with the same vertex set as G , and we add a hyperedge $e = \{u_1, \dots, u_{2\ell}\}$ to G' for all distinct $u_1, \dots, u_{2\ell} \in V$ that induce a 2ℓ -clique in G .
- Run \mathbb{A} on (G', ρ) . Let S be \mathbb{A} 's output.
- Use brute force to find a maximum balanced biclique in S and output it.

Notice that algorithm \mathbb{B} runs in time $(f(2k) + 2^{O(k)}) \cdot n^{O(\ell)} = (f(2k) + 2^{O(k)}) \cdot n^{o(k)}$ as desired, where the term $2^{O(k)}$ comes from the last step.

Next, we claim that, when k is sufficiently large, the output biclique has size at least $t := \lfloor k/8 \rfloor$, which would give us the desired $O(1)$ -approximation ratio. Suppose for the sake of contradiction that this is not true, i.e. that the induced graph $G[S]$ is $K_{t,t}$ -free.

For any sufficiently large k , we have $\ell < t$. Hence, we may apply Lemma 16, which gives the following upper bound on the number of not-necessarily induced copies of $K_{\ell,\ell}$ in $G[S]$:

$$\left(2 \cdot e^{-\frac{\ell^2}{16t}}\right) \cdot \binom{2k}{\ell} \binom{2k - \ell}{\ell} \leq e^{-\Omega(k/g(k)^{0.2})} \cdot \binom{2k}{\ell} \binom{2k - \ell}{\ell}.$$

However, since each hyperedge $e = \{u_1, \dots, u_{2\ell}\}$ corresponds to $\binom{2\ell}{\ell}$ copies of $K_{\ell,\ell}$, the number of hyperedges fully contained in S is thus at most

$$e^{-\Omega(k/g(k)^{0.2})} \cdot \binom{2k}{\ell} \binom{2k - \ell}{\ell} / \binom{2\ell}{\ell} = e^{-\Omega(k/g(k)^{0.2})} \cdot \binom{2k}{2\ell},$$

which is less than $2^{-\rho/g(\rho)} \cdot \binom{2k}{2\ell}$ for any sufficiently large k . This contradicts the approximation guarantee of \mathbb{A} since the optimal solution (i.e. the $2k$ -clique) contains $\binom{2k}{2\ell}$ hyperedges. \blacktriangleleft

References

- 1 Amir Abboud and Greg Bodwin. Reachability preservers: New extremal bounds and approximation algorithms. In *SODA*, pages 1865–1883, 2018. doi:10.1137/1.9781611975031.122.
- 2 Michael Alekhnovich. More on average case vs approximation complexity. *Computational Complexity*, 20(4):755–786, 2011.
- 3 Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri Weinstein. Inapproximability of densest κ -subgraph from average case hardness, 2011.
- 4 Noga Alon, Uriel Feige, Avi Wigderson, and David Zuckerman. Derandomized graph products. *Comput. Complex.*, 5(1):60–75, 1995. doi:10.1007/BF01277956.
- 5 Noga Alon, Raphael Yuster, and Uri Zwick. Finding and counting given length cycles. *Algorithmica*, 17(3):209–223, 1997. doi:10.1007/BF02523189.

- 6 Benny Applebaum. Pseudorandom generators with long stretch and low locality from random local one-way functions. *SIAM J. Comput.*, 42(5):2008–2037, 2013. doi:10.1137/120884857.
- 7 Benny Applebaum, Boaz Barak, and Avi Wigderson. Public-key cryptosystem from different assumptions. In *STOC*, pages 171–180, 2010.
- 8 Sanjeev Arora, Boaz Barak, Markus Brunnermeier, and Rong Ge. Computational complexity and information asymmetry in financial products (extended abstract). In *ICS*, pages 49–65, 2010. URL: <http://conference.iis.tsinghua.edu.cn/ICS2010/content/papers/5.html>.
- 9 Yakov Babichenko, Christos H. Papadimitriou, and Aviad Rubinfeld. Can almost everybody be almost happy? In *ITCS*, pages 1–9, 2016. doi:10.1145/2840728.2840731.
- 10 Maria-Florina Balcan, Christian Borgs, Mark Braverman, Jennifer T. Chayes, and Shang-Hua Teng. Finding endogenously formed communities. In *SODA*, pages 767–783, 2013. doi:10.1137/1.9781611973105.55.
- 11 Boaz Barak, Samuel B. Hopkins, Jonathan A. Kelner, Pravesh K. Kothari, Ankur Moitra, and Aaron Potechin. A nearly tight sum-of-squares lower bound for the planted clique problem. *SIAM J. Comput.*, 48(2):687–735, 2019. doi:10.1137/17M1138236.
- 12 Boaz Barak, Guy Kindler, and David Steurer. On the optimality of semidefinite relaxations for average-case and generalized constraint satisfaction. In *ITCS*, pages 197–214, 2013.
- 13 Siddharth Barman. Approximating Nash equilibria and dense subgraphs via an approximate version of Carathéodory’s theorem. *SIAM J. Comput.*, 47(3):960–981, 2018. doi:10.1137/15M1050574.
- 14 Mihir Bellare, Shafi Goldwasser, Carsten Lund, and A. Russell. Efficient probabilistically checkable proofs and applications to approximations. In *STOC*, pages 294–304, 1993. doi:10.1145/167088.167174.
- 15 Piotr Berman, Arnab Bhattacharyya, Konstantin Makarychev, Sofya Raskhodnikova, and Grigory Yaroslavtsev. Approximation algorithms for spanner problems and directed steiner forest. *Inf. Comput.*, 222:93–107, 2013. doi:10.1016/j.ic.2012.10.007.
- 16 Piotr Berman and Georg Schnitger. On the complexity of approximating the independent set problem. In *STACS*, pages 256–268, 1989. doi:10.1007/BFb0028990.
- 17 Quentin Berthet and Philippe Rigollet. Complexity theoretic lower bounds for sparse principal component detection. In *COLT*, pages 1046–1066, 2013. URL: <http://proceedings.mlr.press/v30/Berthet13.html>.
- 18 Tengyao Wang, Quentin Berthet, and Richard J. Samworth. Statistical and computational trade-offs in estimation of sparse principal components. *The Annals of Statistics*, 2016.
- 19 Amey Bhangale, Rajiv Gandhi, Mohammad Taghi Hajiaghayi, Rohit Khandekar, and Guy Kortsarz. Bicovering: Covering edges with two small subsets of vertices. In *ICALP*, pages 6:1–6:12, 2016. doi:10.4230/LIPIcs.ICALP.2016.6.
- 20 Umang Bhaskar, Yu Cheng, Young Kun Ko, and Chaitanya Swamy. Hardness results for signaling in bayesian zero-sum and network routing games. In *EC*, pages 479–496, 2016. doi:10.1145/2940716.2940753.
- 21 Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. Detecting high log-densities: an $O(n^{1/4})$ approximation for densest k -subgraph. In *STOC*, pages 201–210, 2010. doi:10.1145/1806689.1806719.
- 22 Aditya Bhaskara, Moses Charikar, Aravindan Vijayaraghavan, Venkatesan Guruswami, and Yuan Zhou. Polynomial integrality gaps for strong SDP relaxations of densest k -subgraph. In *SODA*, pages 388–405, 2012. doi:10.1137/1.9781611973099.34.
- 23 Arnab Bhattacharyya, Suprovat Ghoshal, Karthik C. S., and Pasin Manurangsi. Parameterized intractability of even set and shortest vector problem from Gap-ETH. In *ICALP*, pages 17:1–17:15, 2018. doi:10.4230/LIPIcs.ICALP.2018.17.
- 24 Markus Bläser, Balagopal Komarath, and Karteek Sreenivasaiah. Graph pattern polynomials. In *FSTTCS*, pages 18:1–18:13, 2018. doi:10.4230/LIPIcs.FSTTCS.2018.18.
- 25 Mark Braverman, Young Kun-Ko, Aviad Rubinfeld, and Omri Weinstein. ETH hardness for densest- k -subgraph with perfect completeness. In *SODA*, pages 1326–1341, 2017.

10:18 The Strongish Planted Clique Hypothesis and Its Consequences

- 26 Mark Braverman, Young Kun-Ko, and Omri Weinstein. Approximating the best Nash equilibrium in $n^{o(\log n)}$ -time breaks the exponential time hypothesis. In *SODA*, pages 970–982, 2015. doi:10.1137/1.9781611973730.66.
- 27 Matthew Brennan and Guy Bresler. Optimal average-case reductions to sparse PCA: from weak assumptions to strong hardness. In *COLT*, pages 469–470, 2019. URL: <http://proceedings.mlr.press/v99/brennan19b.html>.
- 28 Matthew Brennan, Guy Bresler, Samuel B Hopkins, Jerry Li, and Tselil Schramm. Statistical query algorithms and low-degree tests are almost equivalent. *arXiv preprint*, 2020. arXiv:2009.06107.
- 29 Matthew Brennan, Guy Bresler, and Wasim Huleihel. Reducibility and computational lower bounds for problems with planted sparse structure. In *COLT*, pages 48–166, 2018. URL: <http://proceedings.mlr.press/v75/brennan18a.html>.
- 30 Liming Cai and David W. Juedes. On the existence of subexponential parameterized algorithms. *J. Comput. Syst. Sci.*, 67(4):789–807, 2003. doi:10.1016/S0022-0000(03)00074-6.
- 31 Parinya Chalermsook, Marek Cygan, Guy Kortsarz, Bundit Laekhanukit, Pasin Manurangsi, Danupon Nanongkai, and Luca Trevisan. From Gap-ETH to FPT-Inapproximability: Clique, dominating set, and more. In *FOCS*, pages 743–754, 2017. doi:10.1109/FOCS.2017.74.
- 32 Moses Charikar. Greedy approximation algorithms for finding dense components in a graph. In *APPROX*, pages 84–95, 2000. doi:10.1007/3-540-44436-X_10.
- 33 Moses Charikar, Chandra Chekuri, To-Yat Cheung, Zuo Dai, Ashish Goel, Sudipto Guha, and Ming Li. Approximation algorithms for directed steiner problems. *J. Algorithms*, 33(1):73–91, 1999. doi:10.1006/jagm.1999.1042.
- 34 Moses Charikar, MohammadTaghi Hajiaghayi, and Howard J. Karloff. Improved approximation algorithms for label cover problems. *Algorithmica*, 61(1):190–206, 2011. doi:10.1007/s00453-010-9464-3.
- 35 Chandra Chekuri, Guy Even, Anupam Gupta, and Danny Segev. Set connectivity problems in undirected graphs and the directed steiner network problem. *ACM Trans. Algorithms*, 7(2):18:1–18:17, 2011. doi:10.1145/1921659.1921664.
- 36 Rajesh Chitnis, Andreas Emil Feldmann, and Pasin Manurangsi. Parameterized approximation algorithms for bidirected steiner network problems. In *ESA*, pages 20:1–20:16, 2018. doi:10.4230/LIPIcs.ESA.2018.20.
- 37 Eden Chlamtác, Michael Dinitz, Christian Konrad, Guy Kortsarz, and George Rabanca. The densest k -subhypergraph problem. *SIAM J. Discret. Math.*, 32(2):1458–1477, 2018. doi:10.1137/16M1096402.
- 38 Eden Chlamtác, Michael Dinitz, Guy Kortsarz, and Bundit Laekhanukit. Approximating spanners and directed steiner forest: Upper and lower bounds. In *SODA*, pages 534–553, 2017. doi:10.1137/1.9781611974782.34.
- 39 Eden Chlamtác, Michael Dinitz, and Robert Krauthgamer. Everywhere-sparse spanners via dense subgraphs. In *FOCS*, pages 758–767, 2012. doi:10.1109/FOCS.2012.61.
- 40 Eden Chlamtác, Michael Dinitz, and Yury Makarychev. Minimizing the union: Tight approximations for small set bipartite vertex expansion. In *SODA*, pages 881–899, 2017. doi:10.1137/1.9781611974782.56.
- 41 Eden Chlamtác and Pasin Manurangsi. Sherali-adams integrality gaps matching the log-density threshold. In *APPROX*, pages 10:1–10:19, 2018. doi:10.4230/LIPIcs.APPROX-RANDOM.2018.10.
- 42 Eden Chlamtác, Pasin Manurangsi, Dana Moshkovitz, and Aravindan Vijayaraghavan. Approximation algorithms for label cover and the log-density threshold. In *SODA*, pages 900–919, 2017. doi:10.1137/1.9781611974782.57.
- 43 Radu Curticapean, Holger Dell, and Dániel Marx. Homomorphisms are a good basis for counting small subgraphs. In *STOC*, pages 210–223, 2017. doi:10.1145/3055399.3055502.

- 44 Radu Curticapean and Dániel Marx. Complexity of counting subgraphs: Only the boundedness of the vertex-cover number counts. In *FOCS*, pages 130–139, 2014. doi:10.1109/FOCS.2014.22.
- 45 Marek Cygan, Pawel Komosa, Daniel Lokshtanov, Michal Pilipczuk, Marcin Pilipczuk, and Saket Saurabh. Randomized contractions meet lean decompositions. *CoRR*, abs/1810.06864, 2018. arXiv:1810.06864.
- 46 Marek Cygan, Marcin Pilipczuk, and Michal Pilipczuk. Known algorithms for edge clique cover are probably optimal. *SIAM J. Comput.*, 45(1):67–83, 2016. doi:10.1137/130947076.
- 47 Søren Dahlgaard, Mathias Bæk Tejs Knudsen, and Morten Stöckel. Finding even cycles faster via capped k -walks. In *STOC*, pages 112–120, 2017. doi:10.1145/3055399.3055459.
- 48 Mina Dalirrooyfard, Thuy Duong Vuong, and Virginia Vassilevska Williams. Graph pattern detection: hardness for all induced patterns and faster non-induced cycles. In *STOC*, pages 1167–1178, 2019. doi:10.1145/3313276.3316329.
- 49 Amit Daniely. Complexity theoretic limitations on learning halfspaces. In *STOC*, pages 105–117, 2016. doi:10.1145/2897518.2897520.
- 50 Amit Daniely and Shai Shalev-Shwartz. Complexity theoretic limitations on learning DNF’s. In *COLT*, pages 815–830, 2016. URL: <http://proceedings.mlr.press/v49/daniely16.html>.
- 51 Michael Dinitz, Guy Kortsarz, and Zeev Nutov. Improved approximation algorithm for steiner k -forest with nearly uniform weights. *ACM Trans. Algorithms*, 13(3):40:1–40:16, 2017. doi:10.1145/3077581.
- 52 Irit Dinur and Pasin Manurangsi. ETH-hardness of approximating 2-CSPs and directed steiner network. In *ITCS*, pages 36:1–36:20, 2018. doi:10.4230/LIPIcs.ITCS.2018.36.
- 53 Yevgeniy Dodis and Sanjeev Khanna. Design networks with bounded pairwise distance. In *STOC*, pages 750–759, 1999. doi:10.1145/301250.301447.
- 54 Rodney G. Downey and Michael R. Fellows. Fixed-parameter tractability and completeness II: on completeness for $W[1]$. *Theor. Comput. Sci.*, 141(1&2):109–131, 1995. doi:10.1016/0304-3975(94)00097-3.
- 55 Rodney G. Downey and Michael R. Fellows. *Fundamentals of Parameterized Complexity*. Springer, 2013.
- 56 Paul Erdős and Alfréd Rényi. Asymmetric graphs. *Acta Mathematica Academiae Scientiarum Hungarica*, 14(3-4):295–315, 1963.
- 57 Uriel Feige. Relations between average case complexity and approximation complexity. In *STOC*, pages 534–543, 2002. doi:10.1145/509907.509985.
- 58 Uriel Feige, Guy Kortsarz, and David Peleg. The dense k -subgraph problem. *Algorithmica*, 29(3):410–421, 2001. doi:10.1007/s004530010050.
- 59 Uriel Feige and Robert Krauthgamer. The probable value of the lovász–schrijver relaxations for maximum independent set. *SIAM J. Comput.*, 32(2):345–370, 2003. doi:10.1137/S009753970240118X.
- 60 Uriel Feige and Michael Seltser. On the densest k -subgraph problem. Technical report, Weizmann Institute of Science, Rehovot, Israel, 1997.
- 61 Moran Feldman, Guy Kortsarz, and Zeev Nutov. Improved approximation algorithms for directed steiner forest. *J. Comput. Syst. Sci.*, 78(1):279–292, 2012. doi:10.1016/j.jcss.2011.05.009.
- 62 Vitaly Feldman, Elena Grigorescu, Lev Reyzin, Santosh S. Vempala, and Ying Xiao. Statistical algorithms and a lower bound for detecting planted cliques. *J. ACM*, 64(2):8:1–8:37, 2017. doi:10.1145/3046674.
- 63 Chao Gao, Zongming Ma, and Harrison H Zhou. Sparse CCA: adaptive estimation and computational barriers. *The Annals of Statistics*, 2017.
- 64 Jiong Guo, Rolf Niedermeier, and Ondrej Suchý. Parameterized complexity of arc-weighted directed steiner problems. *SIAM J. Discret. Math.*, 25(2):583–599, 2011. doi:10.1137/100794560.

- 65 Anupam Gupta, Mohammad Taghi Hajiaghayi, Viswanath Nagarajan, and R. Ravi. Dial a ride from k -forest. *ACM Trans. Algorithms*, 6(2):41:1–41:21, 2010. doi:10.1145/1721837.1721857.
- 66 Bruce E. Hajek, Yihong Wu, and Jiaming Xu. Computational lower bounds for community detection on random graphs. In *COLT*, pages 899–928, 2015. URL: <http://proceedings.mlr.press/v40/Hajek15.html>.
- 67 Mohammad Taghi Hajiaghayi and Kamal Jain. The prize-collecting generalized steiner tree problem via a new approach of primal-dual schema. In *SODA*, pages 631–640, 2006. URL: <http://dl.acm.org/citation.cfm?id=1109557.1109626>.
- 68 Elad Hazan and Robert Krauthgamer. How hard is it to approximate the best Nash equilibrium? *SIAM J. Comput.*, 40(1):79–91, 2011. doi:10.1137/090766991.
- 69 Russell Impagliazzo and Ramamohan Paturi. On the complexity of k -SAT. *J. Comput. Syst. Sci.*, 62(2):367–375, 2001. doi:10.1006/jcss.2000.1727.
- 70 Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. Which problems have strongly exponential complexity? *J. Comput. Syst. Sci.*, 63(4):512–530, 2001. doi:10.1006/jcss.2001.1774.
- 71 Mark Jerrum. Large cliques elude the metropolis process. *Random Struct. Algorithms*, 3(4):347–360, 1992. doi:10.1002/rsa.3240030402.
- 72 David S. Johnson. The NP-completeness column: An ongoing guide. *J. Algorithms*, 8(5):438–448, September 1987.
- 73 Richard Karp. Probabilistic analysis of some combinatorial search problems. *Algorithms and Complexity: New Directions and Recent Results*, 1976.
- 74 Richard M. Karp. Reducibility among combinatorial problems. In *Proceedings of a symposium on the Complexity of Computer Computations*, pages 85–103, 1972. doi:10.1007/978-1-4684-2001-2_9.
- 75 Subhash Khot. Ruling out PTAS for graph min-bisection, dense k -subgraph, and bipartite clique. *SIAM J. Comput.*, 36(4):1025–1071, 2006. doi:10.1137/S0097539705447037.
- 76 Guy Kortsarz and David Peleg. On choosing a dense subgraph (extended abstract). In *34th Annual Symposium on Foundations of Computer Science, Palo Alto, California, USA, 3-5 November 1993*, pages 692–701, 1993.
- 77 Tamás Kővári, Vera T Sós, and Pál Turán. On a problem of Zarankiewicz. In *Colloquium Mathematicum*, volume 3, pages 50–57, 1954.
- 78 Mirosław Kowaluk, Andrzej Lingas, and Eva-Marta Lundell. Counting and detecting small subgraphs via equations. *SIAM J. Discret. Math.*, 27(2):892–909, 2013. doi:10.1137/110859798.
- 79 James R. Lee, Prasad Raghavendra, and David Steurer. Lower bounds on the size of semidefinite programming relaxations. In *STOC*, pages 567–576, 2015. doi:10.1145/2746539.2746599.
- 80 Bingkai Lin. The parameterized complexity of k -biclique. In *SODA*, pages 605–615, 2015.
- 81 Andrea Lincoln and Nikhil Vyas. Algorithms and lower bounds for cycles and walks: Small space and sparse graphs. In *ITCS*, pages 11:1–11:17, 2020. doi:10.4230/LIPIcs.ITCS.2020.11.
- 82 Daniel Lokshtanov, Dániel Marx, and Saket Saurabh. Lower bounds based on the exponential time hypothesis. *Bull. EATCS*, 105:41–72, 2011. URL: <http://eatcs.org/beatcs/index.php/beatcs/article/view/92>.
- 83 Daniel Lokshtanov, M. S. Ramanujan, Saket Saurabh, and Meirav Zehavi. Parameterized complexity and approximability of directed odd cycle transversal. *CoRR*, abs/1704.04249, 2017. arXiv:1704.04249.
- 84 Pasin Manurangsi. Almost-polynomial ratio ETH-hardness of approximating densest k -subgraph. In *STOC*, pages 954–961, 2017. doi:10.1145/3055399.3055412.
- 85 Pasin Manurangsi. Inapproximability of maximum edge biclique, maximum balanced biclique and minimum k -cut from the small set expansion hypothesis. In *ICALP*, pages 79:1–79:14, 2017. doi:10.4230/LIPIcs.ICALP.2017.79.
- 86 Pasin Manurangsi and Dana Moshkovitz. Approximating dense max 2-CSPs. In *APPROX*, pages 396–415, 2015. doi:10.4230/LIPIcs.APPROX-RANDOM.2015.396.

- 87 Pasin Manurangsi and Aviad Rubinfeld. Inapproximability of VC dimension and Littlestone's dimension. In *COLT*, pages 1432–1460, 2017. URL: <http://proceedings.mlr.press/v65/manurangsi17a.html>.
- 88 R. Milo, S. Shen-Orr, S. Itzkovitz, N. Kashtan, D. Chklovskii, and U. Alon. Network motifs: Simple building blocks of complex networks. *Science*, 298(5594):824–827, 2002.
- 89 Jaroslav Nesetril and Svatopluk Poljak. On the complexity of the subgraph problem. *Commentationes Mathematicae Universitatis Carolinae*, 26:415–419, 1985.
- 90 Prasad Raghavendra and David Steurer. Graph expansion and the unique games conjecture. In *STOC*, pages 755–764, 2010. doi:10.1145/1806689.1806792.
- 91 Aviad Rubinfeld. Settling the complexity of computing approximate two-player Nash equilibria. In *FOCS*, pages 258–265, 2016. doi:10.1109/FOCS.2016.35.
- 92 Aviad Rubinfeld. Detecting communities is hard (and counting them is even harder). In *ITCS*, pages 42:1–42:13, 2017. doi:10.4230/LIPIcs.ITCS.2017.42.
- 93 Danny Segev and Gil Segev. Approximate k -steiner forests via the lagrangian relaxation technique with internal preprocessing. *Algorithmica*, 56(4):529–549, 2010. doi:10.1007/s00453-008-9189-8.
- 94 Johan Ugander, Lars Backstrom, and Jon M. Kleinberg. Subgraph frequencies: mapping the empirical and extremal geography of large graph collections. In *WWW*, pages 1307–1318, 2013. doi:10.1145/2488388.2488502.
- 95 Virginia Vassilevska Williams. On some fine-grained questions in algorithms and complexity, 2018. ICM survey.
- 96 Virginia Vassilevska Williams and Ryan Williams. Finding, minimizing, and counting weighted subgraphs. *SIAM J. Comput.*, 42(3):831–854, 2013. doi:10.1137/09076619X.
- 97 Raphael Yuster and Uri Zwick. Detecting short directed cycles using rectangular matrix multiplication and dynamic programming. In *SODA*, pages 254–260, 2004. URL: <http://dl.acm.org/citation.cfm?id=982792.982828>.
- 98 David Zuckerman. On unapproximable versions of NP-complete problems. *SIAM Journal on Computing*, 25(6):1293–1304, 1996.