Almost-Polynomial Ratio ETH-Hardness of Approximating Densest k-Subgraph

Pasin Manurangsi[*]
University of California, Berkeley
Berkeley, CA, USA
pasin@berkeley.edu

ABSTRACT

In the Densest k-Subgraph (DkS) problem, given an undirected graph G and an integer k, the goal is to find a subgraph of G on k vertices that contains maximum number of edges. Even though Bhaskara et al.'s state-of-the-art algorithm for the problem achieves only $O(n^{1/4+\varepsilon})$ approximation ratio, previous attempts at proving hardness of approximation, including those under average case assumptions, fail to achieve a polynomial ratio; the best ratios ruled out under any worst case assumption and any average case assumption are only any constant (Raghavendra and Steurer) and $2^{O(\log^{2/3} n)}$ (Alon et al.) respectively.

In this work, we show, assuming the exponential time hypothesis (ETH), that there is no polynomial-time algorithm that approximates DkS to within $n^{1/(\log \log n)^c}$ factor of the optimum, where $c > 0$ is a universal constant independent of n. In addition, our result has *perfect completeness*, meaning that we prove that it is ETH-hard to even distinguish between the case in which G contains a k-clique and the case in which every induced k-subgraph of G has density at most $1/n^{-1/(\log \log n)^c}$ in polynomial time.

Moreover, if we make a stronger assumption that there is some constant $\varepsilon > 0$ such that no subexponential-time algorithm can distinguish between a satisfiable 3SAT formula and one which is only $(1-\varepsilon)$-satisfiable (also known as Gap-ETH), then the ratio above can be improved to $n^{f(n)}$ for any function f whose limit is zero as n goes to infinity (i.e. $f \in o(1)$).

CCS CONCEPTS

• **Theory of computation → Problems, reductions and completeness**;

KEYWORDS

Densest k-Subgraph, Hardness of Approximation

ACM Reference format:
Pasin Manurangsi. 2017. Almost-Polynomial Ratio ETH-Hardness of Approximating Densest k-Subgraph. In *Proceedings of 49th Annual ACM SIGACT Symposium on the Theory of Computing, Montreal, Canada, June 2017 (STOC'17),* 8 pages.
DOI: 10.1145/3055399.3055412

1 INTRODUCTION

In the Densest k-Subgraph (DkS) problem, we are given an undirected graph G on n vertices and a positive integer $k \leqslant n$. The goal is to find a set S of k vertices such that the induced subgraph on S has maximum number of edges. Since the size of S is fixed, the problem can be equivalently stated as finding a k-subgraph (i.e. subgraph on k vertices) with maximum density where density[1] of the subgraph induced on S is $|E(S)|/\binom{|S|}{2}$ and $E(S)$ denotes the set of all edges among the vertices in S.

Densest k-Subgraph, a natural generalization of k-Clique [32], was first formulated and studied by Kortsarz and Peleg [34] in the early 90s. Since then, it has been the subject of intense study in the context of approximation algorithm and hardness of approximation [3, 8, 10, 13–15, 22, 24–27, 33, 41, 46]. Despite this, its approximability still remains wide open and is considered by some to be an important open question in approximation algorithms [13–15].

On the approximation algorithm front, Kortsarz and Peleg [34], in the same work that introduced the problem, gave a polynomial-time $\tilde{O}(n^{0.3885})$-approximation algorithm for DkS. Feige, Kortsarz and Peleg [24] later provided an $O(n^{1/3-\delta})$-approximation for the problem for some constant $\delta \approx 1/60$. This approximation ratio was the best known for almost a decade until Bhaskara et al. [13] invented a log-density based approach which yielded an $O(n^{1/4+\varepsilon})$-approximation for any constant $\varepsilon > 0$. This remains the state-of-the-art approximation algorithm for DkS.

While the above algorithms demonstrate the main progresses of approximations of DkS in general case over the years, many special cases have also been studied. Most relevant to our work is the case where the optimal k-subgraph has high density, in which better approximations are known [10, 26, 37, 47]. The first and most representative algorithm of this kind is that of Feige and Seltser [26], which provides the following guarantee: when the input graph contains a k-clique, the algorithm can find an $(1-\varepsilon)$-dense k-subgraph in $n^{O(\log n/\varepsilon)}$ time. We will refer to this problem of finding densest k-subgraph when the input graph is promised to have a k-clique Densest k-Subgraph *with perfect completeness*.

Although many algorithms have been devised for DkS, relatively little is known regarding its hardness of approximation. While it is commonly believed that the problem is hard to approximate to

[*]This material is based upon work supported by the National Science Foundation under Grants No. CCF 1540685 and CCF 1655215.

STOC'17, Montreal, Canada
© 2017 ACM. 978-1-4503-4528-6/17/06…$15.00
DOI: 10.1145/3055399.3055412

[1] It is worth noting that sometimes density is defined as $|E(S)|/|S|$. For DkS, both definitions of density result in the same objective since $|S| = k$ is fixed. However, our notion is more convenient to deal with as it always lies in $[0, 1]$.

within some polynomial ratio [3, 14], not even a constant factor NP-hardness of approximation is known. To circumvent this, Feige [22] came up with a hypothesis that a random 3SAT formula is hard to refute in polynomial time and proved that, assuming this hypothesis, DkS is hard to approximate to within some constant factor.

Alon et al. [3] later used a similar conjecture regarding random k-AND to rule out polynomial-time algorithms for DkS with any constant approximation ratio. Moreover, they proved hardnesses of approximation of DkS under the following *Planted Clique Hypothesis* [31, 35]: there is no polynomial-time algorithm that can distinguish between a typical Erdős-Rényi random graph $\mathcal{G}(n, 1/2)$ and one in which a clique of size polynomial in n (e.g. $n^{1/3}$) is planted. Assuming this hypothesis, Alon et al. proved that no polynomial-time algorithm approximates DkS to within any constant factor. They also showed that, when the hypothesis is strengthened to rule out not only polynomial-time but also super-polynomial time algorithms for the Planted Clique problem, their inapproximability guarantee for DkS can be improved. In particular, if no $n^{O(\sqrt{\log n})}$-time algorithm solves the Planted Clique problem, then $2^{O(\log^{2/3} n)}$-approximation for DkS cannot be achieved in polynomial time.

There are also several inapproximability results of DkS based on worst-case assumptions. Khot [33] showed, assuming NP $\not\subseteq$ BPTIME($2^{n^{\varepsilon}}$) for some constant $\varepsilon > 0$, that no polynomial-time algorithm can approximate DkS to within $(1 + \delta)$ factor where $\delta > 0$ is a constant depending only on ε; the proof is based on a construction of a "quasi-random" PCP, which is then used in place of a random 3SAT in a reduction similar to that from [22].

While no inapproximability of DkS is known under the Unique Games Conjecture, Raghavendra and Steurer [41] showed that a strengthened version of it, in which the constraint graph is required to satisfy a "small-set expansion" property, implies that DkS is hard to approximate to within any constant ratio.

Recently, Braverman et al. [15], showed, under the exponential time hypothesis (ETH), which will be stated shortly, that, for some constant $\varepsilon > 0$, no $n^{\tilde{O}(\log n)}$-time algorithm can approximate Densest k-Subgraph with perfect completeness to within $(1 + \varepsilon)$ factor. It is worth noting here that their result matches almost exactly with the previously mentioned Feige-Seltser algorithm [26].

Since none of these inapproximability results achieve a polynomial ratio, there have been efforts to prove better lower bounds for more restricted classes of algorithms. For example, Bhaskara et al. [14] provided polynomial ratio lower bounds against strong SDP relaxations of DkS. Specifically, for the Sum-of-Squares hierarchy, they showed integrality gaps of $n^{2/53-\varepsilon}$ and n^{ε} against $n^{\Omega(\varepsilon)}$ and $n^{1-O(\varepsilon)}$ levels of the hierarchy respectively. (See also [18, 36] in which 2/53 in the exponent was improved to 1/14.) Unfortunately, it is unlikely that these lower bounds can be translated to inapproximability results and the question of whether any polynomial-time algorithm can achieve subpolynomial approximation ratio for DkS remains an intriguing open question.

1.1 Our Results

In this work, we rule out, under the exponential time hypothesis (i.e. no subexponential time algorithm can solve 3SAT; see Hypothesis 2.1), polynomial-time approximation algorithms for DkS (even with perfect completeness) with slightly subpolynomial ratio:

Theorem 1.1. *There is a constant $c > 0$ such that, assuming ETH, no polynomial-time algorithm can, given a graph G on n vertices and a positive integer $k \leqslant n$, distinguish between the following two cases:*

- *There exist k vertices of G that induce a k-clique.*
- *Every k-subgraph of G has density at most $n^{-1/(\log\log n)^c}$.*

If we assume a stronger assumption that it takes exponential time to even distinguish between a satisfiable 3SAT formula and one which is only $(1 - \varepsilon)$-satisfiable for some constant $\varepsilon > 0$ (aka Gap-ETH; see Hypothesis 2.2), the ratio can be improved to $n^{f(n)}$ for any $f \in o(1)$:

Theorem 1.2. *For every function $f \in o(1)$, assuming Gap-ETH, no polynomial-time algorithm can, given a graph G on n vertices and a positive integer $k \leqslant n$, distinguish between the following two cases:*

- *There exist k vertices of G that induce a k-clique.*
- *Every k-subgraph of G has density at most $n^{-f(n)}$.*

We remark that, for DkS with perfect completeness, the aforementioned Feige-Seltser algorithm achieves an n^{ε}-approximation in time $n^{O(1/\varepsilon)}$ for every $\varepsilon > 0$ [26]. Hence, the ratios in our theorems cannot be improved to some fixed polynomial and the ratio in Theorem 1.2 is tight in this sense.

Comparison to Previous Results. In terms of inapproximability ratios, the ratios ruled out in this work are almost polynomial and provides a vast improvement over previous results. Prior to our result, the best known ratio ruled out under any worst case assumption is only any constant [41] and the best ratio ruled out under any average case assumption is only $2^{O(\log^{2/3} n)}$ [3]. In addition, our results also have perfect completeness, which was only achieved in [15] under ETH and in [3] under the Planted Clique Hypothesis but not in [22, 33, 41].

Regarding the assumptions our results are based upon, the average case assumptions used in [3, 22] are incomparable to ours. The assumption NP $\not\subseteq$ BPTIME($2^{n^{\varepsilon}}$) used in [33] is also incomparable to ours since, while not stated explicitly, ETH and Gap-ETH by default focus only on deterministic algorithms and our reductions are also deterministic. The strengthened Unique Games Conjecture used in [41] is again incomparable to ours as it is a statement that a specific problem is NP-hard. Finally, although Braverman et al.'s result [15] also relies on ETH, its relation to our results is more subtle. Specifically, their reduction time is only $2^{\tilde{\Theta}(\sqrt{m})}$ where m is the number of clauses, meaning that they only need to assume that 3SAT $\notin$ DTIME($2^{\tilde{\Theta}(\sqrt{m})}$) to rule out a constant ratio polynomial-time approximation for DkS. However, as we will see in Theorem 3.1, even to achieve a constant gap, our reduction time is $2^{\tilde{\Omega}(m^{3/4})}$. Hence, if 3SAT somehow ends up in DTIME($2^{\tilde{\Theta}(m^{3/4})}$)\DTIME($2^{\tilde{\Theta}(\sqrt{m})}$), their result will still hold whereas ours will not even imply constant ratio inapproximability for DkS.

Implications of Our Results. One of the reasons that DkS has received significant attention in the approximation algorithm community is due to its connections to many other problems. Most relevant to our work are the problems to which there are reductions from DkS that preserve approximation ratios to within some polynomial[2]. These problems include Densest At-Most-k-Subgraph [5],

[2] These are problems whose $O(\rho)$-approximation gives an $O(\rho^c)$-approximation for DkS for some constant c.

Smallest *m*-Edge Subgraph [17], Steiner *k*-Forest [28] and Quadratic Knapsack [40]. For brevity, we do not define these problems here. We refer interested readers to cited sources for their definitions and reductions from D*k*S to respective problems. We also note that this list is by no means exhaustive and there are indeed numerous other problems with similar known connections to D*k*S. Our results also imply hardness of approximation results with similar ratios to D*k*S for such problems:

Corollary 1.3. *For some constant $c > 0$, assuming ETH, there is no polynomial-time $n^{1/(\log \log n)^c}$-approximation algorithm for* Densest At-Most-*k*-Subgraph, Smallest *m*-Edge Subgraph, Steiner *k*-Forest, Quadratic Knapsack. *Moreover, for any function $f \in o(1)$, there is no polynomial-time $n^{f(n)}$-approximation algorithm for any of these problems, unless Gap-ETH is false.*

2 PRELIMINARIES AND NOTATIONS

We use $\exp(x)$ and $\log(x)$ to denote e^x and $\log_2(x)$ respectively. polylog n is used as a shorthand for $O(\log^c n)$ for some absolute constant c. For any set S, $\mathscr{P}(S) := \{T \mid T \subseteq S\}$ denotes the power set of S. For any non-negative integer $t \leqslant |S|$, we use $\binom{S}{t} := \{T \in \mathscr{P}(S) \mid |T| = t\}$ to denote the collection of all subsets of S of size t.

Throughout this work, we only concern with simple unweighted undirected graphs. Recall that the density of a graph $G = (V, E)$ on $N \geqslant 2$ vertices is $|E|/\binom{N}{2}$. We say that a graph is *α-dense* if its density is α. Moreover, for every $t \in \mathbb{N}$, we view each element of V^t as a t-size ordered multiset of V. $(L, R) \in V^t \times V^t$ is said to be a *labelled copy of a t-biclique* (or $K_{t,t}$) in G if, for every $u \in L$ and $v \in R$, $u \neq v$ and $(u, v) \in E$. The number of labelled copies of $K_{t,t}$ in G is the number of all such (L, R)'s.

2.1 Exponential Time Hypotheses

One of our results is based on the exponential time hypothesis (ETH), a conjecture proposed by Impagliazzo and Paturi [29] which asserts that 3SAT cannot be solved in subexponential time:

Hypothesis 2.1 (ETH [29]). *No $2^{o(m)}$-time algorithm can decide whether any 3SAT formula with m clauses*[3] *is satisfiable.*

Another hypothesis used in this work is Gap-ETH, a strengthened version of the ETH, which essentially states that even approximating 3SAT to some constant ratio takes exponential time:

Hypothesis 2.2 (Gap-ETH [21, 38]). *There exists a constant $\varepsilon > 0$ such that no $2^{o(m)}$-time algorithm can, given a 3SAT formula ϕ with m clauses*[4]*, distinguish between the case where ϕ is satisfiable and the case where $\mathrm{val}(\phi) \leqslant 1 - \varepsilon$. Here $\mathrm{val}(\phi)$ denote the maximum fraction of clauses of ϕ satisfied by any assignment.*

2.2 Nearly-Linear Size PCPs and Subexponential Time Reductions

The celebrated PCP Theorem [6, 7], which lies at the heart of virtually all known NP-hardness of approximation results, can be viewed as a polynomial-time reduction from 3SAT to a gap version of 3SAT, as stated below. While this perspective is a rather narrow viewpoint of the theorem that leaves out the fascinating relations between parameters of PCPs, it will be the most convenient for our purpose.

Theorem 2.3 (PCP Theorem [6, 7]). *For some constant $\varepsilon > 0$, there exists a polynomial-time reduction that takes a 3SAT formula φ and produces a 3SAT formula ϕ such that*

- (Completeness) *if φ is satisfiable, then ϕ is satisfiable, and,*
- (Soundness) *if φ is unsatisfiable, then* $\mathrm{val}(\phi) \leqslant 1 - \varepsilon$.

Following the first proofs of the PCP Theorem, considerable efforts have been made to improve the trade-offs between the parameters in the theorem. One such direction is to try to reduce the size of the PCP, which, in the above formulation, translates to reducing the size of ϕ relative to φ. On this front, it is known that the size of ϕ can be made nearly-linear in the size of φ [12, 20, 39]. For our purpose, we will use Dinur's PCP Theorem [20], which has a blow-up of only polylogarithmic in the size of ϕ:

Theorem 2.4 (Dinur's PCP Theorem [20]). *For some constant $\varepsilon, d > 0$, there exists a polynomial-time reduction that takes a 3SAT formula φ with m clauses and produces another 3SAT formula ϕ with $m' = O(m \operatorname{polylog} m)$ clauses such that*

- (Completeness) *if φ is satisfiable, then ϕ is satisfiable, and,*
- (Soundness) *if φ is unsatisfiable, then* $\mathrm{val}(\phi) \leqslant 1 - \varepsilon$*, and,*
- (Bounded Degree) *each variable of ϕ appears in $\leqslant d$ clauses.*

Note that Dinur's PCP, combined with ETH, implies a lower bound of $2^{\Omega(m/\operatorname{polylog} m)}$ on the running time of algorithms that solve the gap version of 3SAT, which is only a factor of $O(\operatorname{polylog} m)$ in the exponent off from Gap-ETH. Putting it differently, Gap-ETH is closely related to the question of whether a linear size PCP, one where the size blow-up is only constant instead of polylogarithmic, exists; its existence would mean that Gap-ETH is implied by ETH.

Under the exponential time hypothesis, nearly-linear size PCPs allow us to start with an instance ϕ of the gap version of 3SAT and reduce, in subexponential time, to another problem. As long as the time spent in the reduction is $2^{o(m/\operatorname{polylog} m)}$, we arrive at a lower bound for the problem. Arguably, Aaronson et al. [1] popularized this method, under the name *birthday repetition*, by using such a reduction of size $2^{\tilde{\Omega}(\sqrt{m})}$ to prove ETH-hardness for free games and dense CSPs. Without going into any detail now, let us mention that the name birthday repetition comes from the use of the birthday paradox in their proof and, since its publication, their work has inspired many inapproximability results [9, 15, 16, 19, 38, 43, 45]. Our result too is inspired by this line of work and, as we will see soon, part of our proof also contains a birthday-type paradox.

3 THE REDUCTION AND PROOFS OF THE MAIN THEOREMS

The reduction from the gap version of 3SAT to D*k*S is simple. Given a 3SAT formula ϕ on n variables $x_1, \ldots, x_n$ and an integer $1 \leqslant \ell \leqslant n$, we construct a graph[5] $G_{\phi,\ell} = (V_{\phi,\ell}, E_{\phi,\ell})$ as follows:

[3] In its original form, the running time lower bound is exponential in the number of variables not the number of clauses; however, thanks to the sparsification lemma of Impagliazzo et al. [30], both versions are equivalent.

[4] As noted by Dinur [21], a subsampling argument can be used to make the number of clauses linear in the number of variables, meaning that the conjecture remains the same even when m denotes the number of variables.

[5] For interested readers, we note that our graph is not the same as the FGLSS graph [23] of the PCP in which the verifier reads ℓ random variables and accepts if no clause is violated; while this graph has the same vertex set as ours, the edges are different since we check that no clause between the two vertices is violated, which is not checked in

- Its vertex set $V_{\phi,\ell}$ contains all partial assignments to ℓ variables, i.e., each vertex is $\{(x_{i_1}, b_{i_1}), \dots, (x_{i_\ell}, b_{i_\ell})\}$ where $x_{i_1}, \dots, x_{i_\ell}$ are ℓ distinct variables and $b_{i_1}, \dots, b_{i_\ell} \in \{0, 1\}$ are the bits assigned to them.
- We connect two vertices $\{(x_{i_1}, b_{i_1}), \dots, (x_{i_\ell}, b_{i_\ell})\}$ and $\{(x_{i'_1}, b_{i'_1}), \dots, (x_{i'_\ell}, b_{i'_\ell})\}$ by an edge iff the two partial assignments are consistent (i.e. no variable is assigned 0 in one vertex and 1 in another), and, every clause in ϕ all of whose variables are from $x_{i_1}, \dots, x_{i_\ell}, x_{i'_1}, \dots, x_{i'_\ell}$ is satisfied by the partial assignment induced by the two vertices.

Clearly, if $\text{val}(\phi) = 1$, the $\binom{n}{\ell}$ vertices corresponding to a satisfying assignment induce a clique. Our main technical contribution is proving that, when $\text{val}(\phi) \leqslant 1 - \varepsilon$, every $\binom{n}{\ell}$-subgraph is sparse:

Theorem 3.1. *For any $d, \varepsilon > 0$, there exists $\delta > 0$ such that, for any 3SAT formula ϕ on n variables such that* $\text{val}(\phi) \leqslant 1 - \varepsilon$ *and each variable appears in at most d clauses and for any integer $\ell \in [n^{3/4}/\delta, n/2]$, any $\binom{n}{\ell}$-subgraph of $G_{\phi,\ell}$ has density $\leqslant 2^{-\delta \ell^4/n^3}$.*

We remark that there is nothing special about 3SAT; we can start with any boolean CSP and end up with a similar result, albeit the soundness deteriorates as the arity of the CSP grows. However, it is crucial that the variables are boolean; in fact, Braverman et al. [15] considered a graph similar to ours for 2CSPs but they were unable to achieve subconstant soundness since their variables were not boolean[6]. In particular, there is a non-boolean 2CSP with low value which results in the graph having a biclique of size larger than $\binom{n}{\ell}$ (see Appendix A), i.e., one cannot get an inapproximability ratio more than two starting from a non-boolean CSP.

Once we have Theorem 3.1, the inapproximability results of DkS (Theorem 1.1 and 1.2) can be easily proved by applying the theorem with appropriate choices of ℓ. We defer these proofs to Subsection 3.2. For now, let us turn our attention to the proof of Theorem 3.1. To prove the theorem, we resort to the following lemma due to Alon [2], which states that every dense graph contains many labelled copies of bicliques:

Lemma 3.2 ([2, Lemma 2.1]). *Any α-dense graph G on $N \geqslant 2$ vertices has at least $(\alpha/2)^{t^2} N^{2t}$ labelled copies of $K_{t,t}$ for all $t \in \mathbb{N}$.*

Equipped with Lemma 3.2, our proof strategy is to bound the number of labelled copies of $K_{t,t}$ in $G_{\phi,\ell}$ where t is to be chosen later. To argue this, we will need some additional notations:

- First, let $A_\phi := \{(x_1, 0), (x_1, 1), \dots, (x_n, 0), (x_n, 1)\}$ be the set of all single-variable partial assignments. Observe that $V_{\phi,\ell} \subseteq \binom{A_\phi}{\ell}$, i.e., each $u \in V_{\phi,\ell}$ is a subset of A_ϕ of size ℓ.
- Let $\mathcal{A} : (V_{\phi,\ell})^t \to \mathscr{P}(A_\phi)$ be a "flattening" function that, on input $T \in (V_{\phi,\ell})^t$, outputs the set of all single-variable partial assignments that appear in at least one vertex in T. In other words, when each vertex u is viewed as a subset of A_ϕ, we can write $\mathcal{A}(T)$ simply as $\bigcup_{u \in T} u$.
- Let $\mathcal{K}_{t,t} := \{(L, R) \in (V_{\phi,\ell})^t \times (V_{\phi,\ell})^t \mid \forall u \in L, \forall v \in R, u \neq v \wedge (u, v) \in E_{\phi,\ell}\}$ denote the set of all labelled copies of $K_{t,t}$ in $G_{\phi,\ell}$ and, for each $A, B \subseteq A_\phi$, let $\mathcal{K}_{t,t}(A, B) := \{(L, R) \in \mathcal{K}_{t,t} \mid \mathcal{A}(L) = A, \mathcal{A}(R) = B\}$ denote the set of all $(L, R) \in \mathcal{K}_{t,t}$ with $\mathcal{A}(L) = A$ and $\mathcal{A}(R) = B$.

The number of labelled copies of $K_{t,t}$ in $G_{\phi,\ell}$ can be written as

$$|\mathcal{K}_{t,t}| = \sum_{A, B \subseteq A_\phi} |\mathcal{K}_{t,t}(A, B)|. \tag{1}$$

To bound $|\mathcal{K}_{t,t}|$, we will prove the following bound on $|\mathcal{K}_{t,t}(A, B)|$.

Lemma 3.3. *Let ϕ, n, ℓ, d and ε be as in Theorem 3.1. There exists $\lambda > 0$ depending only on d and ε such that, for any $t \in \mathbb{N}$ and any $A, B \subseteq A_\phi$, $|\mathcal{K}_{t,t}(A, B)| \leqslant \left(2^{-\lambda \ell^2/n} \binom{n}{\ell}\right)^{2t}$.*

Before we prove the above lemma, let us see how Lemma 3.2 and Lemma 3.3 imply Theorem 3.1.

Proof of Theorem 3.1. Assume w.l.o.g. that $\lambda \leqslant 1$. Pick $\delta = \lambda^2/8$ and $t = (4/\lambda)(n^2/\ell^2)$. From Lemma 3.3 and (1), we have

$$|\mathcal{K}_{t,t}| \leqslant 2^{4n} \cdot \left(2^{-\lambda \ell^2/n} \binom{n}{\ell}\right)^{2t} \leqslant (2^{-\lambda \ell^2/n})^t \cdot \binom{n}{\ell}^{2t}$$

where the second inequality comes from our choice of t; note that t is chosen so that the 2^{4n} factor is consumed by $2^{-\lambda \ell^2/n}$ from Lemma 3.3. Finally, consider any $\binom{n}{\ell}$-subgraph of $G_{\phi,\ell}$. By the above bound, it contains at most $(2^{-\lambda \ell^2/n})^t \cdot \binom{n}{\ell}^{2t}$ labelled copies of $K_{t,t}$. Thus, from Lemma 3.2 and from $\ell \geqslant n^{3/4}/\delta$, its density is at most $2 \cdot 2^{-\lambda \ell^2/(nt)} = 2 \cdot 2^{-2\delta \ell^4/n^3} \leqslant 2^{-\delta \ell^4/n^3}$ as desired. □

We now move on to the proof of Lemma 3.3.

Proof of Lemma 3.3. First, notice that if (x, b) appears in A and $(x, \neg b)$ appears in B for some variable x and bit b, then $\mathcal{K}_{t,t}(A, B) = \emptyset$; this is because, for any L with $\mathcal{A}(L) = A$ and R with $\mathcal{A}(R) = B$, there exist $u \in L$ and $v \in R$ that contain (x, b) and $(x, \neg b)$ respectively, meaning that there is no edge between u and v and, thus, $(L, R) \notin \mathcal{K}_{t,t}(A, B)$. Hence, from now on, we can assume that, if (x, b) appears in one of A, B, then the other does not contain $(x, \neg b)$. Observe that this implies that, for each variable x, its assignments can appear in A and B at most two times[7] in total. This in turn implies that $|A| + |B| \leqslant 2n$.

Let us now argue that $|\mathcal{K}_{t,t}(A, B)| \leqslant \binom{n}{\ell}^{2t}$; while this is not the bound we are looking for yet, it will serve as a basis for our argument later. For every $(L, R) \in \mathcal{K}_{t,t}(A, B)$, observe that, since $\mathcal{A}(L) = A$ and $\mathcal{A}(R) = B$, we have $L \in \binom{A}{\ell}^t$ and $R \in \binom{B}{\ell}^t$. This implies that $\mathcal{K}_{t,t}(A, B) \subseteq \binom{A}{\ell}^t \times \binom{B}{\ell}^t$. Hence,

$$|\mathcal{K}_{t,t}(A, B)| \leqslant \binom{|A|}{\ell}^t \binom{|B|}{\ell}^t. \tag{2}$$

the FGLSS graph. It is possible to modify our proof to make it work for this FGLSS graph. However, the soundness guarantee for the FGLSS graph is worse.

[6] Any satisfiable boolean 2CSP is solvable in polynomial time so one cannot start with a boolean 2CSP either.

[7] This is where we use the fact that the variables are boolean. For non-boolean CSPs, each variable x can appear more than two times in one of A or B alone, which can indeed be problematic (see Appendix A).

Moreover, $\binom{|A|}{\ell}\binom{|B|}{\ell}$ can be further bounded as

$$\begin{aligned}\binom{|A|}{\ell}\binom{|B|}{\ell} &= \frac{1}{(\ell!)^2}\prod_{i=0}^{\ell-1}(|A|-i)(|B|-i)\\ \text{(From AM-GM Inequality)} &\leqslant \frac{1}{(\ell!)^2}\prod_{i=0}^{\ell-1}\left(\frac{|A|+|B|}{2}-i\right)^2\\ \text{(From from } |A|+|B|\leqslant 2n) &\leqslant \binom{n}{\ell}^2 \qquad (3)\end{aligned}$$

Combining (2) and (3) indeed yields $|\mathcal{K}_{t,t}(A,B)| \leqslant \binom{n}{\ell}^{2t}$.

Inequality (2) is very crude; we include all elements of $\binom{A}{\ell}$ and $\binom{B}{\ell}$ as candidates for vertices in L and R respectively. However, as we will see soon, only tiny fraction of elements of $\binom{A}{\ell}, \binom{B}{\ell}$ can actually appear in L, R when $(L,R) \in \mathcal{K}_{t,t}(A,B)$. To argue this, let us categorize the variables into three groups:

- x is *terrible* iff its assignments appear at most once in total in A and B (i.e. $|\{(x,0),(x,1)\}\cap A| + |\{(x,0),(x,1)\}\cap B| \leqslant 1$).
- x is *good* iff, for some $b \in \{0,1\}$, $(x,b) \in A \cap B$. Note that this implies that $(x,\neg b) \notin A \cup B$.
- x is *bad* iff either $\{(x,0),(x,1)\} \subseteq A$ or $\{(x,0),(x,1)\} \subseteq B$.

The next and last step of the proof is where birthday-type paradoxes come in. Before we continue, let us briefly demonstrate the ideas behind this step by considering the following extreme cases:

- If all variables are terrible, then $|A| + |B| \leqslant n$ and (3) can be immediately tightened.
- If all variables are bad, assume w.l.o.g. that, for at least half of variables x's, $\{(x,0),(x,1)\} \subseteq A$. Consider a random element u of $\binom{A}{\ell}$. Since u is a set of random ℓ distinct elements of A, there will, in expectation, be $\Omega(\ell^2/n)$ variables x's with $(x,0),(x,1) \in u$. However, the presence of such x's means that u is not a valid vertex. Moreover, it is not hard to turn this into the following probabilistic statement: with probability at most $2^{-\Omega(\ell^2/n)}$, u contains at most one of $(x,0),(x,1)$ for every variable x. In other words, only $2^{-\Omega(\ell^2/n)}$ fraction of elements of $\binom{A}{\ell}$ are valid vertices, which yields the desired bound on $|\mathcal{K}_{t,t}(A,B)|$.
- If all variables are good, then $A = B$ is simply an assignment to all the variables. Since $\text{val}(\phi) \leqslant 1-\varepsilon$, at least εm clauses are unsatisfied by this assignment. As we will argue below, every element of $\binom{A}{\ell}$ that contains two variables from some unsatisfied clause cannot be in L for any $(L,R) \in \mathcal{K}_{t,t}(A,B)$. This means that there are $\Theta_\varepsilon(m) \geqslant \Omega_\varepsilon(n)$ prohibited pairs of variables that cannot appear together. Again, similar to the previous case, it is not hard to argue that only $2^{-\Omega_{\varepsilon,d}(\ell^2/n)}$ fraction of elements of $\binom{A}{\ell}$ can be candidates for vertices of L.

To turn this intuition into a bound on $|\mathcal{K}_{t,t}(A,B)|$, we need the following inequality. Its proof is straightforward and is deferred to Subsection 3.1.

Proposition 3.4. *Let U be any set and $P \subseteq \binom{U}{2}$ be any set of pairs of elements of U such that each element of U appears in at most q pairs. For any positive integer $2 \leqslant r \leqslant |U|$, the probability that a random element of $\binom{U}{r}$ does not contain both elements of any pair in P is at most* $\exp\left(-\frac{|P|r^2}{4q|U|^2}\right)$.

We are now ready to formalize the above intuition and finish the proof of Lemma 3.3. For the sake of convenience, denote the sets of good, bad and terrible variables by X_g, X_b and X_t respectively. Moreover, let $\beta := \varepsilon/(100d)$ and pick $\lambda = \min\{-\log(1-\beta/2), \beta/64, \varepsilon/(384d)\}$. To refine the bound on the size of $\mathcal{K}_{t,t}(A,B)$, consider the following three cases:

(1) $|X_t| \geqslant \beta n$. Since each $x \in X_t$ contributes at most one to $|A|+|B|$, $|A|+|B| \leqslant (1-\beta/2)(2n)$. Hence, we can improve (3) to $\binom{|A|}{\ell}\binom{|B|}{\ell} \leqslant \binom{(1-\beta/2)n}{\ell}^2$. Thus, $|\mathcal{K}_{t,t}(A,B)|$ is bounded above by

$$\binom{(1-\beta/2)n}{\ell}^{2t} \leqslant \left((1-\beta/2)^\ell\binom{n}{\ell}\right)^{2t} \leqslant \left(2^{-\lambda\ell^2/n}\binom{n}{\ell}\right)^{2t}$$

where the last inequality comes from $\lambda \leqslant -\log(1-\beta/2)$ and $\ell > \ell^2/n$.

(2) $|X_b| \geqslant \beta n$. Since each $x \in X_b$ appears either in A or B, one of A and B must contain assignments to at least $(\beta/2)n$ variables in X_b. Assume w.l.o.g. that A satisfies this property. Let X_b^L be the set of all $x \in X_b$ whose assignments appear in A. Note that $|X_b^L| \geqslant (\beta/2)n$.

Observe that an element $u \in \binom{A}{\ell}$ is not a valid vertex if it contains both $(x,0)$ and $(x,1)$ for some $x \in X_b^L$. We invoke Proposition 3.4 with $U = A$, $P = \{\{(x,0),(x,1)\} \mid x \in X_b^L\}$, $q = 1$ and $r = \ell$, which implies that a random element of $\binom{A}{\ell}$ does not contain any prohibited pairs in P with probability at most $\exp\left(-\frac{|X_b^L|\ell^2}{4|A|^2}\right) \leqslant \exp\left(-\frac{(\beta/2)n\ell^2}{4(2n)^2}\right)$, which is at most $2^{-2\lambda\ell^2/n}$ because $\lambda \leqslant \beta/64$. In other words, at most $2^{-2\lambda\ell^2/n}$ fraction of elements of $\binom{A}{\ell}$ are valid vertices. This gives us the following upper bound on $|\mathcal{K}_{t,t}(A,B)|$:

$$\left(2^{-2\lambda\ell^2/n}\cdot\binom{|A|}{\ell}\right)^t\cdot\binom{|B|}{\ell}^t \overset{(3)}{\leqslant} \left(2^{-\lambda\ell^2/n}\binom{n}{\ell}\right)^{2t}$$

(3) $|X_t| < \beta n$ and $|X_b| < \beta n$. In this case, $|X_g| > (1-2\beta)n$. Let S denote the set of clauses whose variables all lie in X_g. Since each variable appears in at most d clauses, $|S| > m-(2\beta n)d \geqslant (1-\varepsilon/2)m$ where the second inequality comes from our choice of β and from $m \geqslant n/3$.

Consider the partial assignment $f : X_g \to \{0,1\}$ induced by A and B, i.e., $f(x) = b$ iff $(x,b) \in A, B$. Since $\text{val}(\phi) \leqslant 1-\varepsilon$, the number of clauses in S satisfied by f is at most $(1-\varepsilon)m$. Hence, at least $\varepsilon m/2$ clauses in S are unsatisfied by f. Denote the set of such clauses by S_{UNSAT}.

Fix a clause $C \in S_{\text{UNSAT}}$ and let x, y be two different variables in C. We claim that x, y cannot appear together in any vertex of L for any $(L,R) \in \mathcal{K}_{t,t}(A,B)$. Suppose for the sake of contradiction that $(x, f(x))$ and $(y, f(y))$ both appear in $u \in L$ for some $(L,R) \in \mathcal{K}_{t,t}(A,B)$. Let $z \in X_g$ be another variable[8] in C. Since $(z, f(z)) \in B$, some vertex

[8] If C contains two variables, let $z = x$. Note that we can assume w.l.o.g. that C contains at least two variables.

$v \in R$ contains $(z, f(z))$. Thus, there is no edge between u and v in $G_{\phi,\ell}$, which contradicts with $(L, R) \in \mathcal{K}_{t,t}$.

We can now appeal to Proposition 3.4 with $U = A$, $q = 2d$, $r = \ell$ and P be the prohibited pairs described above. This implies that with probability at most $\exp\left(-\frac{|P|\ell^2}{8d|A|^2}\right) \leqslant \exp\left(-\frac{\varepsilon\ell^2}{192dn}\right)$, a random element of $\binom{A}{\ell}$ contains no prohibited pair from P. In other words, at most $\exp\left(-\frac{\varepsilon\ell^2}{192dn}\right)$ fraction of elements of $\binom{A}{\ell}$ can be candidates for each element of L for $(L, R) \in \mathcal{K}_{t,t}(A, B)$. This gives the following sharpened upper bound on $|\mathcal{K}_{t,t}(A, B)|$:

$$\left(\exp\left(-\frac{\varepsilon\ell^2}{192dn}\right) \cdot \binom{|A|}{\ell}\right)^t \cdot \binom{|B|}{\ell}^t \overset{(3)}{\leqslant} \left(2^{-\varepsilon\ell^2/(384dn)} \cdot \binom{n}{\ell}\right)^{2t}.$$

Since we picked $\lambda \leqslant \varepsilon/(384d)$, $|\mathcal{K}_{t,t}(A, B)|$ is once again bounded above by $\left(2^{-\lambda\ell^2/n}\binom{n}{\ell}\right)^{2t}$.

In all three cases, we have $|\mathcal{K}_{t,t}(A, B)| \leqslant \left(2^{-\lambda\ell^2/n}\binom{n}{\ell}\right)^{2t}$, completing the proof of Lemma 3.3. □

3.1 Proof of Proposition 3.4

Proof of Proposition 3.4. We first construct $P' \subseteq P$ such that each element of U appears in at most one pair in P' as follows. Start out by marking every pair in P as active and, as long as there are active pairs left, include one in P' and mark every pair that shares an element of U with this pair as inactive. Since each element of U appears in at most q pairs in P, we mark at most $2q$ pairs as inactive per each inclusion. This implies that $|P'| \geqslant |P|/(2q)$.

Suppose that $P' = \{\{a_1, b_1\}, \dots, \{a_{|P'|}, b_{|P'|}\}\}$. Let u be a random element of $\binom{U}{r}$. For each $i = 1, \dots, |P'|$, we have

$$\Pr[\{a_i, b_i\} \not\subseteq u] = 1 - \frac{\binom{|U|-2}{r-2}}{\binom{|U|}{r}} \leqslant 1 - \frac{r^2}{2|U|^2} \leqslant \exp\left(-\frac{r^2}{2|U|^2}\right).$$

If u does not contain both elements of any pairs in P, it does not contain both elements of any pairs in P'. The probability of the latter can be written as

$$\Pr\left[\bigwedge_{i=1}^{|P'|}\{a_i, b_i\} \not\subseteq u\right] = \prod_{i=1}^{|P'|}\Pr\left[\{a_i, b_i\} \not\subseteq u \,\middle|\, \bigwedge_{j=1}^{i-1}\{a_j, b_j\} \not\subseteq u\right].$$

In addition, since $a_1, b_1, \dots, a_{|P'|}, b_{|P'|}$ are distinct, it is not hard to see that $\Pr\left[\{a_i, b_i\} \not\subseteq u \middle| \bigwedge_{j=1}^{i-1}\{a_j, b_j\} \not\subseteq u\right] \leqslant \Pr[\{a_i, b_i\} \not\subseteq u]$. Hence, we can upper bound $\Pr\left[\bigwedge_{i=1}^{|P'|}\{a_i, b_i\} \not\subseteq u\right]$ by

$$\prod_{i=1}^{|P'|}\Pr[\{a_i, b_i\} \not\subseteq u] \leqslant \left(\exp\left(-\frac{r^2}{2|U|^2}\right)\right)^{|P'|} \leqslant \exp\left(-\frac{|P|r^2}{4q|U|^2}\right),$$

completing the proof of Proposition 3.4. □

3.2 Proofs of Inapproximability Results of DkS

In this subsection, we prove Theorem 1.1 and 1.2. The proof of Theorem 1.1 is simply by combining Dinur's PCP Theorem and Theorem 3.1 with $\ell = m/\operatorname{polylog} m$, as stated below.

Proof of Theorem 1.1. For any 3SAT formula φ with m clauses, use Theorem 2.4 to produce ϕ with $m' = O(m \operatorname{polylog} m)$ clauses such that each variable appears in at most d clauses. Let ζ be a constant such that $m' = O(m \log^{\zeta} m)$ and let $\ell = m/\log^2 m$. Let us consider the graph $G_{\phi,\ell}$ with $k = \binom{n}{\ell}$ where n is the number of variables of ϕ. Let N be the number of vertices of $G_{\phi,\ell}$. Observe that $N = 2^{\ell}\binom{n}{\ell} \leqslant n^{2\ell} \leqslant (m')^{O(\ell)} = 2^{O(\ell \log m')} = 2^{o(m)}$.

If φ is satisfiable, ϕ is also satisfiable and it is obvious that $G_{\phi,\ell}$ contains an induced k-clique. Otherwise, If φ is unsatisfiable, $\mathrm{val}(\phi) \leqslant 1 - \varepsilon$. From Theorem 3.1, any k-subgraph of $G_{\phi,\ell}$ has density at most $2^{-\Omega(\ell^4/n^3)} \leqslant 2^{-\Omega(m/\log^{3\zeta+8} m)} = N^{-\Omega(1/(\log\log N)^{3\zeta+8})}$, which is at most $N^{-1/(\log\log N)^{3\zeta+9}}$ when m is sufficiently large. Hence, if there is a polynomial-time algorithm that can distinguish between the two cases in Theorem 1.1 when $c = 3\zeta + 9$, then there also exists an algorithm that solves 3SAT in time $2^{o(m)}$, contradicting with ETH. □

The proof of Theorem 1.2 is even simpler since, under Gap-ETH, we have the gap version of 3SAT to begin with. Hence, we can directly apply Theorem 3.1 without going through Dinur's PCP:

Proof of Theorem 1.2. Let ϕ be any 3SAT formula with m clauses such that each variable appears in $O(1)$ clauses[9]. Let $\ell = m\sqrt[5]{f(m)}$ and consider the graph $G_{\phi,\ell}$ with $k = \binom{n}{\ell}$ where n is the number of variables of ϕ. The number of vertices N of $G_{\phi,\ell}$ is $2^{\ell}\binom{n}{\ell} \leqslant 2^{\ell}\left(\frac{en}{\ell}\right)^{\ell} \leqslant 2^{O(\ell \log(m/\ell))} = 2^{O(m\sqrt[5]{f(m)}\log(1/f(m)))} = 2^{o(m)}$ where the last inequality follows from $f \in o(1)$.

The completeness is again obvious. For the soundness, if $\mathrm{val}(\phi) \leqslant 1 - \varepsilon$, from Theorem 3.1, any k-subgraph of $G_{\phi,\ell}$ has density at most $2^{-\Omega(\ell^4/n^3)} \leqslant 2^{-\Omega(mf(m)^{4/5})} \leqslant N^{-\Omega(f(m)^{4/5})}$, which is at most[10] $N^{-f(N)}$ when m is sufficiently large. Hence, if there is a polynomial-time algorithm that can distinguish between the two cases in Theorem 1.2, then there also exists an algorithm that solves the gap version of 3SAT in time $2^{o(m)}$, contradicting with Gap-ETH. □

4 CONCLUSION AND OPEN QUESTIONS

In this work, we provide a subexponential time reduction from the gap version of 3SAT to DkS and prove that it establishes an almost-polynomial ratio hardness of approximation of the latter under ETH and Gap-ETH. Even with our results, however, approximability of DkS still remains wide open. Namely, it is still not known whether it is NP-hard to approximate DkS to within some constant factor, and, no polynomial ratio hardness of approximation is yet known.

Although our results appear to almost resolve the second question, it still seems out of reach with our current knowledge of hardness of approximation. In particular, to achieve a polynomial ratio hardness for DkS, it is plausible that one has to prove a long-standing conjecture called *the sliding scale conjecture (SSC)* [11]. In short, SSC essentially states that Label Cover, a problem used as starting points of almost all NP-hardness of approximation results, is NP-hard to approximate to within some polynomial ratio. Note here that polynomial ratio hardness for Label Cover is not even

[9] We can assume w.l.o.g. that each variable appears in at most $O(1)$ clauses [38, p.21].

[10] Assume w.l.o.g. that f is decreasing; otherwise take $f'(m) = \sup_{m' \geqslant m} f(m')$ instead.

known under stronger assumptions such as ETH or Gap-ETH; we refer the readers to [21] for more detailed discussions on the topic.

Apart from the approximability of D*k*S, our results also prompt the following question: since previous techniques, such as Feige's Random 3SAT Hypothesis [22], Khot's Quasi-Random PCP [33], the Small Set Expansion Conjecture [41] and the Planted Clique Hypothesis [31, 35], that were successful in showing inapproximability of D*k*S also gave rise to hardnesses of approximation of many problems that are not known to be APX-hard such as Sparsest Cut and Min Bisection [4, 42], is it possible to modify our construction to prove inapproximability for these problems as well?

A A COUNTEREXAMPLE TO OBTAINING A SUBCONSTANT SOUNDNESS FROM NON-BOOLEAN CSPS

Here we sketch an example due to Rubinstein [44] of a non-boolean 2CSP ϕ with low value for which the graph $G_{\phi,\ell}$ contains a large biclique. For a non-boolean 2CSP, we define the graph $G_{\phi,\ell}$ similar to that of a 3SAT formula except that now the vertices contains all $\{(x_{i_1},\sigma_{i_1}),\dots,(x_{i_\ell},\sigma_{i_\ell})\}$ for all distinct variables $x_{i_1},\dots,x_{i_\ell}$ and all $\sigma_{i_1},\dots,\sigma_{i_\ell}\in\Sigma$ where Σ is the alphabet of the CSP.

Consider any non-boolean 2CSP instance ϕ on variables $x_1,\dots,x_n$ such that there is no constraint between $X_1 := \{x_1,\dots,x_{n/2}\}$ and $X_2 = \{x_{n/2+1},\dots,x_n\}$ and each variable appears in $\leqslant d$ constraints. Let L the set of all vertices u such that every variable in u belongs to X_1 and no constraint is contained in u. Define R similarly for X_2. Clearly, (L,R) forms a biclique and it is not hard to see that $|L|,|R| \geqslant |\Sigma|^\ell\binom{n/2-(d+1)\ell}{\ell}$. Since $|\Sigma|\geqslant 3$, this value is $\geqslant\binom{n}{\ell}$ for all $\ell\leqslant\frac{n}{6(d+2)}$. Hence, for such ℓ, $G_{\phi,\ell}$ contains a biclique of size $\binom{n}{\ell}$.

Finally, note that there are several ways to define constraints within X_1 and X_2 so that val(ϕ) is bounded away from one. For instance, we can make each side a random 2-XOR formula, which results in val$(\phi)\leqslant 1/2+O(1/\sqrt{d})$ w.h.p. Thus, if we start from a non-boolean CSP, the largest gap we can hope to get is only two.

Note that the instance above is rather extreme as it consists of two disconnected components. Hence, it is still possible that, if the starting CSP has more specific properties (e.g. expanding constraint graph), then one can arrive at a gap of more than two.

ACKNOWLEDGEMENT

I am grateful to Aviad Rubinstein, Prasad Raghavendra and Luca Trevisan for invaluable discussions throughout various stages of this work. Without them, this work would not have been possible. Furthermore, I thank Daniel Reichman and Igor Shinkar for dicussions on a related problem which inspire part of the proof. I also thank Igor for pointing me to [2]. Finally, I thank anonymous reviewers for useful comments on an earlier draft of this work.

REFERENCES

[1] Scott Aaronson, Russell Impagliazzo, and Dana Moshkovitz. 2014. AM with Multiple Merlins. In *IEEE CCC 2014*. 44–55.

[2] Noga Alon. 2002. Testing subgraphs in large graphs. *Random Struct. Algorithms* 21, 3-4 (2002), 359–370.

[3] Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri Weinstein. 2011. Inapproximabilty of Densest k-Subgraph from Average Case Hardness. (2011). Unpublished Manuscript.

[4] Christoph Ambuhl, Monaldo Mastrolilli, and Ola Svensson. 2007. Inapproximability Results for Sparsest Cut, Optimal Linear Arrangement, and Precedence Constrained Scheduling. In *IEEE FOCS 2007*. 329–337.

[5] Reid Andersen and Kumar Chellapilla. 2009. Finding Dense Subgraphs with Size Bounds. In *WAW 2009*. 25–37.

[6] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. 1998. Proof Verification and the Hardness of Approximation Problems. *J. ACM* 45, 3 (May 1998), 501–555.

[7] Sanjeev Arora and Shmuel Safra. 1998. Probabilistic Checking of Proofs: A New Characterization of NP. *J. ACM* 45, 1 (Jan. 1998), 70–122.

[8] Yuichi Asahiro, Refael Hassin, and Kazuo Iwama. 2002. Complexity of finding dense subgraphs. *Discrete Appl. Math.* 121, 1âĂŞ3 (2002), 15 – 26.

[9] Yakov Babichenko, Christos H. Papadimitriou, and Aviad Rubinstein. 2016. Can Almost Everybody be Almost Happy?. In *ACM ITCS 2016*. 1–9.

[10] Siddharth Barman. 2015. Approximating Nash Equilibria and Dense Bipartite Subgraphs via an Approximate Version of Caratheodory's Theorem. In *ACM STOC 2015*. 361–369.

[11] Mihir Bellare, Shafi Goldwasser, Carsten Lund, and Alexander Russell. 1993. Efficient Probabilistically Checkable Proofs and Applications to Approximations. In *ACM STOC 1993*. 294–304.

[12] Eli Ben-Sasson and Madhu Sudan. 2008. Short PCPs with Polylog Query Complexity. *SIAM J. Comput.* 38, 2 (May 2008), 551–607.

[13] Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. 2010. Detecting high log-densities: an $O(n^{1/4})$ approximation for densest k-subgraph. In *ACM STOC 2010*. 201–210.

[14] Aditya Bhaskara, Moses Charikar, Aravindan Vijayaraghavan, Venkatesan Guruswami, and Yuan Zhou. 2012. Polynomial Integrality Gaps for Strong SDP Relaxations of Densest k-subgraph. In *ACM-SIAM SODA 2012*. 388–405.

[15] Mark Braverman, Young Kun-Ko, Aviad Rubinstein, and Omri Weinstein. 2017. ETH Hardness for Densest-k-Subgraph with Perfect Completeness. In *ACM-SIAM SODA 2017*. 1326–1341.

[16] Mark Braverman, Young Kun-Ko, and Omri Weinstein. 2015. Approximating the best Nash Equilibrium in $n^{o(\log n)}$-time breaks the Exponential Time Hypothesis. In *ACM-SIAM SODA 2015*. 970–982.

[17] Eden Chlamtac, Michael Dinitz, and Robert Krauthgamer. 2012. Everywhere-Sparse Spanners via Dense Subgraphs. In *IEEE FOCS 2012*. 758–767.

[18] Eden Chlamtác, Pasin Manurangsi, Dana Moshkovitz, and Aravindan Vijayaraghavan. 2017. Approximation Algorithms for Label Cover and The Log-Density Threshold. In *ACM-SIAM SODA 2017*. 900–919.

[19] Argyrios Deligkas, John Fearnley, and Rahul Savani. 2016. Inapproximability Results for Approximate Nash Equilibria. *CoRR* abs/1608.03574 (2016).

[20] Irit Dinur. 2007. The PCP Theorem by Gap Amplification. *J. ACM* 54, 3, Article 12 (June 2007).

[21] Irit Dinur. 2016. Mildly exponential reduction from gap 3SAT to polynomial-gap label-cover. *ECCC* 23 (2016), 128.

[22] Uriel Feige. 2002. Relations Between Average Case Complexity and Approximation Complexity. In *ACM STOC 2002*. 534–543.

[23] Uriel Feige, Shafi Goldwasser, László Lovász, Shmuel Safra, and Mario Szegedy. 1996. Interactive Proofs and the Hardness of Approximating Cliques. *J. ACM* 43, 2 (1996), 268–292.

[24] Uriel Feige, Guy Kortsarz, and David Peleg. 2001. The Dense k-Subgraph Problem. *Algorithmica* 29, 3 (2001), 410–421.

[25] Uriel Feige and Michael Langberg. 2001. Approximation Algorithms for Maximization Problems Arising in Graph Partitioning. *J. Algorithms* 41, 2 (Nov. 2001), 174–211.

[26] Uriel Feige and Michael Seltser. 1997. *On the densest k-subgraph problem.* Technical Report. Weizmann Institute of Science, Rehovot, Israel.

[27] Doron Goldstein and Michael Langberg. 2009. The Dense k Subgraph problem. *CoRR* abs/0912.5327 (2009).

[28] Mohammad Taghi Hajiaghayi and Kamal Jain. 2006. The Prize-collecting Generalized Steiner Tree Problem via a New Approach of Primal-dual Schema. In *ACM-SIAM SODA 2006*. 631–640.

[29] Russell Impagliazzo and Ramamohan Paturi. 2001. On the Complexity of k-SAT. *J. Comput. Syst. Sci.* 62, 2 (March 2001), 367–375.

[30] Russell Impagliazzo, Ramamohan Paturi, and Francis Zane. 2001. Which Problems Have Strongly Exponential Complexity? *J. Comput. Syst. Sci.* 63, 4 (Dec. 2001), 512–530.

[31] Mark Jerrum. 1992. Large Cliques Elude the Metropolis Process. *Random Struct. Algorithms* 3, 4 (1992), 347–359.

[32] Richard M. Karp. 1972. Reducibility Among Combinatorial Problems. In *Proceedings of a symposium on the Complexity of Computer Computations*. 85–103.

[33] Subhash Khot. 2006. Ruling Out PTAS for Graph Min-Bisection, Dense k-Subgraph, and Bipartite Clique. *SIAM J. Comput.* 36, 4 (2006), 1025–1071.

[34] Guy Kortsarz and David Peleg. 1993. On Choosing a Dense Subgraph (Extended Abstract). In *IEEE FOCS 1993*. 692–701.

[35] Luděk Kučera. 1995. Expected Complexity of Graph Partitioning Problems. *Discrete Appl. Math.* 57, 2-3 (Feb. 1995), 193–212.

[36] Pasin Manurangsi. 2015. *On Approximating Projection Games*. Master's thesis. Massachusetts Institute of Technology, Cambridge, MA.

[37] Pasin Manurangsi and Dana Moshkovitz. 2015. Approximating Dense Max 2-CSPs. In *APPROX 2015*. 396–415.

[38] Pasin Manurangsi and Prasad Raghavendra. 2016. A Birthday Repetition Theorem and Complexity of Approximating Dense CSPs. *CoRR* aabs/1607.02986 (2016).

[39] Dana Moshkovitz and Ran Raz. 2008. Two-query PCP with Subconstant Error. *J. ACM* 57, 5, Article 29 (June 2008).

[40] David Pisinger. 2007. The quadratic knapsack problemâĂŤa survey. *Discrete Appl. Math.* 155, 5 (2007), 623 – 648.

[41] Prasad Raghavendra and David Steurer. 2010. Graph Expansion and the Unique Games Conjecture. In *ACM STOC 2010*. 755–764.

[42] Prasad Raghavendra, David Steurer, and Madhur Tulsiani. 2012. Reductions between Expansion Problems. In *IEEE CCC 2012*. 64–73.

[43] Aviad Rubinstein. 2015. ETH-Hardness for Signaling in Symmetric Zero-Sum Games. *CoRR* abs/1510.04991 (2015).

[44] Aviad Rubinstein. 2016. Personal communication. (2016).

[45] Aviad Rubinstein. 2016. Settling the Complexity of Computing Approximate Two-Player Nash Equilibria. In *IEEE FOCS 2016*. 258–265.

[46] Anand Srivastav and Katja Wolf. 1998. Finding Dense Subgraphs with Semidefinite Programming. In *APPROX 1998*. 181–191.

[47] Akiko Suzuki and Takeshi Tokuyama. 2008. Dense Subgraph Problems with Output-density Conditions. *ACM Trans. Algorithms* 4, 4, Article 43 (Aug. 2008).