

A New Conjecture on Hardness of 2-CSP's with Implications to Hardness of Densest k -Subgraph and Other Problems

Julia Chuzhoy

Toyota Technological Institute at Chicago, USA. Email: cjulia@ttic.edu.

Mina Dalirrooyfard

Massachusetts Institute of Technology, USA. Email: minad@mit.edu.

Vadim Grinberg

Weizmann Institute of Science, Israel. Email: vadim.grinberg@weizmann.ac.il.

Zihan Tan

DIMACS, Rutgers University, USA. Email: zihantan1993@gmail.com.

Abstract

We propose a new conjecture on hardness of 2-CSP's, and show that new hardness of approximation results for Densest k -Subgraph and several other problems, including a graph partitioning problem, and a variation of the Graph Crossing Number problem, follow from this conjecture. The conjecture can be viewed as occupying a middle ground between the d -to-1 conjecture, and hardness results for 2-CSP's that can be obtained via standard techniques, such as Parallel Repetition combined with standard 2-prover protocols for the 3SAT problem. We hope that this work will motivate further exploration of hardness of 2-CSP's in the regimes arising from the conjecture. We believe that a positive resolution of the conjecture will provide a good starting point for other hardness of approximation proofs.

Another contribution of our work is proving that the problems that we consider are roughly equivalent from the approximation perspective. Some of these problems arose in previous work, from which it appeared that they may be related to each other. We formalize this relationship in this work.

2012 ACM Subject Classification Mathematics of computing \rightarrow Approximation algorithms

Keywords and phrases Hardness of Approximation, Densest k -Subgraph

Digital Object Identifier 10.4230/LIPIcs.ITCS.2023.17

Related Version *Full Version*: <https://arxiv.org/abs/2211.05906>

Funding *Julia Chuzhoy*: Supported in part by NSF grant CCF-2006464.

Mina Dalirrooyfard: Part of the work done at Toyota Technological Institute at Chicago.

Vadim Grinberg: Part of the work done at Toyota Technological Institute at Chicago. Supported in part by NSF grant CCF-2006464.

Zihan Tan: Supported by a grant to DIMACS from the Simons Foundation (820931) and NSF grant CCF-2006464.

1 Introduction

In this paper we consider several graph optimization problems, the most prominent and extensively studied of which is Densest k -Subgraph. One of the main motivations of this work is to advance our understanding of the approximability of these problems. Towards this goal, we propose a new conjecture on the hardness of a class of 2-CSP problems, and we show that new hardness of approximation results for all these problems follow from this conjecture. We believe that the conjecture is interesting in its own right, as it can be



© Julia Chuzhoy, Mina Dalirrooyfard, Vadim Grinberg, and Zihan Tan; licensed under Creative Commons License CC-BY 4.0

14th Innovations in Theoretical Computer Science Conference (ITCS 2023).

Editor: Yael Tauman Kalai; Article No. 17; pp. 17:1–17:24

Leibniz International Proceedings in Informatics



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

43 seen as occupying a middle ground between the d -to-1 conjecture, and the type of hardness
 44 of approximation results that one can obtain for 2-CSP problems via standard methods
 45 (such as using constant-factor hardness of approximation results for 3-SAT, combined with
 46 standard 2-prover protocols and Parallel Repetition). While our conditional hardness of
 47 approximation proofs are combinatorial and algorithmic in nature, we hope that this work
 48 will inspire complexity theorists to study the conjecture, and also lead to other hardness of
 49 approximation proofs that combine both combinatorial and algebraic techniques.

50 We prove a new conditional hardness of approximation result for Densest k -Subgraph
 51 based on our conjecture. In addition to the Densest k -Subgraph problem, we study three other
 52 problems. The first problem, called (r,h)-Graph Partitioning, recently arose in the hardness
 53 of approximation proof of the Node-Disjoint Paths problem of [18], who mention that the
 54 problem appears similar to Densest k -Subgraph, but could not formalize this intuition. We
 55 also study a new problem that we call Dense k -Coloring, that can be viewed as a natural
 56 middle ground between Densest k -Subgraph and (r,h)-Graph Partitioning. The fourth problem
 57 that we study is a variation of the notoriously difficult Minimum Crossing Number problem,
 58 that we call Maximum Bounded-Crossing Subgraph. This problem also arose implicitly in [18].
 59 We show that all four problems are roughly equivalent from the approximation perspective, in
 60 the regime where the approximation factors are somewhat large (but some of our reductions
 61 require quasi-polynomial time). We then derive conditional hardness of approximation results
 62 for all these problems based on these reductions and the conditional hardness of Densest
 63 k -Subgraph.

64 The main contribution of this paper is thus twofold: first, we propose a new conjecture on
 65 hardness of CSP's and show that a number of interesting hardness of approximation results
 66 follow from it. Second, we establish a close connection between the four problems that we
 67 study. The remainder of the Introduction is organized as follows. We start by providing a
 68 brief overview of the four problems that we study in this paper. We then state our conjecture
 69 on hardness of CSP's and put it into context with existing results and well-known conjectures.
 70 Finally, we provide a more detailed overview of our results and techniques.

71 Densest k -Subgraph.

72 In the Densest k -Subgraph problem, given an n -vertex graph G and an integer $k > 1$, the
 73 goal is to compute a subset S of k vertices of G , while maximizing the number of edges in
 74 $G[S]$. Densest k -Subgraph is one of the most basic graph optimization problems that has
 75 been studied extensively (see e.g. [2, 7–10, 12, 15, 25–30, 36, 41, 44, 46, 47, 52]). At the same
 76 time it seems notoriously difficult, and despite this extensive work, our understanding of its
 77 approximability is still incomplete. The best current approximation algorithm for Densest
 78 k -Subgraph, due to [8], achieves, for every $\varepsilon > 0$, an $O(n^{1/4+\varepsilon})$ -approximation, in time
 79 $n^{O(1/\varepsilon)}$. Even though the problem appears to be very hard, its hardness of approximation
 80 proof has been elusive. For example, no constant-factor hardness of approximation proofs
 81 for Densest k -Subgraph are currently known under the standard $P \neq NP$ assumption, or
 82 even the stronger assumption that $NP \not\subseteq \text{BPTIME}(n^{\text{poly} \log n})$. In a breakthrough result,
 83 Khot [36] proved a factor- c hardness of approximation for Densest k -Subgraph, for some
 84 small constant c , assuming that $NP \not\subseteq \cap_{\varepsilon>0} \text{BPTIME}(2^{n^\varepsilon})$. Several other papers proved
 85 constant and super-constant hardness of approximation results for Densest k -Subgraph under
 86 *average-case* complexity assumptions: namely that no efficient algorithm can refute random
 87 3-SAT or random k -AND formulas [2, 25]. Additionally, a factor $2^{\Omega(\log^{2/3} n)}$ -hardness of
 88 approximation was shown under assumptions on solving Planted Clique [2]. In a recent
 89 breakthrough, Manurangsi [46] proved that, under the Exponential Time Hypothesis (ETH),

90 the Densest k -Subgraph problem is hard to approximate to within factor $n^{1/(\log \log n)^c}$, for
 91 some constant c . Proving a super-constant hardness of Densest k -Subgraph under weaker
 92 complexity assumptions remains a tantalizing open question that we attempt to address
 93 in this paper. Unfortunately, it seems unlikely that the techniques of [46] can yield such a
 94 result. In this paper we show that, assuming the conjecture on hardness of 2-CSP that we
 95 introduce, Densest k -Subgraph is NP-hard to approximate to within factor $2^{(\log n)^\varepsilon}$, for some
 96 constant $\varepsilon > 0$.

97 The (r, h) -Graph Partitioning Problem.

98 A recent paper [18] on the hardness of approximation of the Node-Disjoint Paths (NDP)
 99 problem formulated and studied a new graph partitioning problem, called (r, h) -Graph Par-
 100 titioning. The input to the problem is a graph G , and two integers, r and h . The goal
 101 is to compute r vertex-disjoint subgraphs H_1, \dots, H_r of G , such that for each $1 \leq i \leq r$,
 102 $|E(H_i)| \leq h$, while maximizing $\sum_{i=1}^r |E(H_i)|$. A convenient intuitive way of thinking about
 103 this problem is that we are interested in obtaining a balanced partition of the graph G into
 104 r vertex-disjoint subgraphs, so that the subgraphs contain sufficiently many edges. Unlike
 105 standard graph partitioning problems, that typically aim to minimize the number of edges
 106 connecting the different subgraphs in the solution, our goal is to maximize the total number
 107 of edges that are contained in the subgraphs. In order to avoid trivial solutions, in which
 108 one of the subgraphs contains almost the entire graph G , and the remaining subgraphs are
 109 almost empty, we place an upper bound h on the number of edges that each subgraph may
 110 contribute towards the solution. Note that the subgraphs H_i of G in the solution need not
 111 be vertex-induced subgraphs.

112 The work of [18] attempted to use (r, h) -Graph Partitioning as a proxy problem for proving
 113 hardness of approximation of NDP. Their results imply that NDP is at least as hard to
 114 approximate as (r, h) -Graph Partitioning, to within polylogarithmic factors. In order to prove
 115 hardness of NDP, it would then be sufficient to show that (r, h) -Graph Partitioning is hard
 116 to approximate. Unfortunately, [18] were unable to do so. Instead, they considered a
 117 generalization of (r, h) -Graph Partitioning, called (r, h) -Graph Partitioning with Bundles. They
 118 showed that NDP is at least as hard as (r, h) -Graph Partitioning with Bundles, and then proved
 119 hardness of this new problem. In the (r, h) -Graph Partitioning with Bundles problem, the input
 120 is the same as in (r, h) -Graph Partitioning, but now graph G must be bipartite, and, for every
 121 vertex v , we are given a partition $\mathcal{B}(v)$ of the set of edges incident to v into subsets that are
 122 called *bundles*. We require that, in a solution (H_1, \dots, H_r) to the problem, for every vertex
 123 $v \in V(G)$, and every bundle $\beta \in \mathcal{B}(v)$, at most one edge of β contributes to the solution;
 124 in other words, at most one edge of β may lie in $\bigcup_i E(H_i)$. This is a somewhat artificial
 125 problem, but this definition allows one to bypass some of the barriers that arise when trying
 126 to prove hardness of (r, h) -Graph Partitioning from existing hardness results for CSP's.

127 It was noted in [18] that the (r, h) -Graph Partitioning problem resembles the Densest
 128 k -Subgraph problem for two reasons. First, in Densest k -Subgraph, the goal is to compute a
 129 dense subgraph of a given graph, with a prescribed number of vertices. One can think of
 130 (r, h) -Graph Partitioning as the problem of computing many vertex-disjoint dense subgraphs
 131 of a given graph. Second, natural hardness of approximation proofs for both problems
 132 seem to run into the same barriers. It is therefore natural to ask: (i) Can we prove that
 133 the (r, h) -Graph Partitioning problem itself is hard to approximate? In particular, can the
 134 techniques of [18] be exploited in order to obtain such a proof? and (ii) Can we formalize
 135 this intuitive connection between (r, h) -Graph Partitioning and Densest k -Subgraph? In this
 136 paper we make progress on both these questions. Our conditional hardness result for

137 Densest k -Subgraph indeed builds on the ideas from [18] for proving hardness of (r,h)-Graph
 138 Partitioning with Bundles. We also provide “almost” approximation-preserving reductions
 139 between (r,h)-Graph Partitioning and Densest k -Subgraph: we show that, if there is an efficient
 140 factor $\alpha(n)$ -approximation algorithm for Densest k -Subgraph, then there is a randomized
 141 efficient factor $O(\alpha(n^2) \cdot \text{poly log } n)$ -approximation algorithm to (r,h)-Graph Partitioning. We
 142 also provide a reduction in the opposite direction: we prove that, if there is an efficient $\alpha(n)$ -
 143 approximation algorithm for (r,h)-Graph Partitioning, then there is a randomized algorithm for
 144 Densest k -Subgraph, that achieves approximation factor $O((\alpha(n^{O(\log n)}))^3 \cdot \log^2 n)$, in time
 145 $n^{O(\log n)}$. Therefore, we prove that Densest k -Subgraph and (r,h)-Graph Partitioning are roughly
 146 equivalent from the approximation perspective (at least for large approximation factors and
 147 quasi-polynomial running times). Combined with our conditional hardness of approximation
 148 for Densest k -Subgraph, our results show that, assuming the conjecture on hardness of 2-CSP
 149 that we introduce, for some constant $0 < \varepsilon \leq 1/2$, there is no efficient $2^{(\log n)^\varepsilon}$ -approximation
 150 algorithm for (r,h)-Graph Partitioning, unless $\text{NP} \subseteq \text{BPTIME}(n^{O(\log n)})$.

151 Maximum Bounded-Crossing Subgraph.

152 The third problem that we study is a variation of the classical Minimum Crossing Number
 153 problem. In the Minimum Crossing Number problem, given an input n -vertex graph G , the
 154 goal is to compute a drawing of G in the plane while minimizing the number of crossings in the
 155 drawing. We define the notions of graph drawing and crossings formally in the Preliminaries,
 156 but these notions are quite intuitive and the specifics of the definition are not important in
 157 this high-level overview.

158 The Minimum Crossing Number problem was initially introduced by Turán [54] in 1944,
 159 and has been extensively studied since then (see, e.g., [13, 14, 16, 19, 20, 32, 33], and also [48–51]
 160 for excellent surveys). But despite all this work, most aspects of the problem are still
 161 poorly understood. A long line of work [16, 17, 20, 21, 24, 32, 33, 43] has recently led to
 162 the first sub-polynomial approximation algorithm for the problem in *low degree graphs*.
 163 Specifically, [21] obtain a factor $O(2^{O((\log n)^{7/8} \log \log n)} \cdot \Delta^{O(1)})$ -approximation algorithm
 164 for Minimum Crossing Number, where Δ is the maximum vertex degree. To the best of
 165 our knowledge, no non-trivial approximation algorithms are known for the problem when
 166 vertex degrees in the input graph G can be arbitrary. However, on the negative side, only
 167 APX-hardness is known for the problem [3, 11]. As the current understanding of the Minimum
 168 Crossing Number problem from the approximation perspective is extremely poor, it is natural
 169 to study hardness of approximation of its variants.

170 Let us consider two extreme variations of the Minimum Crossing Number problem. The
 171 first variant is the Minimum Crossing Number problem itself, where we need to draw an input
 172 graph G in the plane with fewest crossings. The second variant is where we need to compute
 173 a subgraph G' of the input graph G that is planar, while maximizing $|E(G')|$. The latter
 174 problem has a simple constant-factor approximation algorithm, obtained by letting G' be
 175 any spanning forest of G (this is since a planar n -vertex graph may only have $O(n)$ edges).

176 In this paper we study a variation of the Minimum Crossing Number problem, that we
 177 call Maximum Bounded-Crossing Subgraph, which can be viewed as an intermediate problem
 178 between these two extremes. In the Maximum Bounded-Crossing Subgraph problem, given
 179 an n -vertex graph G and an integer $L > 0$, the goal is to compute a subgraph $H \subseteq G$,
 180 such that H has a plane drawing with at most L crossings, while maximizing $|E(H)|$.
 181 Unless we are interested in constant approximation factors, this problem is only interesting
 182 when the bound L on the number of crossings is $\Omega(n)$. This is since, from the Crossing
 183 Number Inequality [1, 42], if $|E(G)| \geq 4|V(G)|$, then the crossing number of G is at least

184 $\Omega(|E(G)|^3/|V(G)|^2)$. Therefore, for $L = O(n)$, a spanning tree provides a constant-factor
 185 approximation to the problem. We emphasize that the focus here is on dense graphs, whose
 186 crossing number may be as large as $\Omega(n^4)$.

187 The Maximum Bounded-Crossing Subgraph problem was implicitly used in [18] for proving
 188 hardness of approximation of NDP, as an intermediate problem, in the reduction from
 189 (r,h)-Graph Partitioning with Bundles to NDP. Their work suggests that there may be a
 190 connection between (r,h)-Graph Partitioning and Maximum Bounded-Crossing Subgraph, even
 191 though the two problems appear quite different. In this paper we prove that the two
 192 problems are roughly equivalent from the approximation perspective: if there is an efficient
 193 factor $\alpha(n)$ -approximation algorithm for (r,h)-Graph Partitioning, then there is an efficient
 194 $O(\alpha(n) \cdot \text{poly log } n)$ -approximation algorithm for Maximum Bounded-Crossing Subgraph. On
 195 the other hand, an efficient $\alpha(n)$ -approximation algorithm for Maximum Bounded-Crossing
 196 Subgraph implies an efficient $O((\alpha(n))^2 \cdot \text{poly log } n)$ -approximation algorithm for (r,h)-Graph
 197 Partitioning. Combined with our conditional hardness of approximation for (r,h)-Graph
 198 Partitioning, we get that, assuming the conjecture on hardness of 2-CSP that we introduce,
 199 for some constant $0 < \varepsilon \leq 1/2$ there is no efficient $2^{(\log n)^\varepsilon}$ -approximation algorithm for
 200 Maximum Bounded-Crossing Subgraph, unless $\text{NP} \subseteq \text{BPTIME}(n^{O(\log n)})$.

201 Dense k -Coloring.

202 The fourth and last problem that we consider is Dense k -Coloring. In this problem, the input
 203 is an n -vertex graph G and an integer k , such that n is an integral multiple of k . The goal
 204 is to partition $V(G)$ into n/k disjoint subsets $S_1, \dots, S_{n/k}$, of cardinality k each, so as to
 205 maximize $\sum_{i=1}^{n/k} |E(S_i)|$. This problem can be viewed as an intermediate problem between
 206 Densest k -Subgraph and (r,h)-Graph Partitioning. The connection to (r,h)-Graph Partitioning
 207 seems clear: in both problems, the goal is to compute a large collection of disjoint subgraphs
 208 of the input graph G , that contain many edges of G . While in (r,h)-Graph Partitioning we
 209 place a limit on the number of edges in each subgraph, in Dense k -Coloring we require that
 210 each subgraph contains exactly k vertices. The connection to the Densest k -Subgraph problem
 211 is also clear: while in Densest k -Subgraph the goal is to compute a single dense subgraph
 212 of G containing k vertices, in Dense k -Coloring we need to partition G into many dense
 213 subgraphs, containing k vertices each. We show reductions between the Dense k -Coloring
 214 and the Densest k -Subgraph problem in both directions, that provide very similar guarantees
 215 to the reductions between (r,h)-Graph Partitioning and Densest k -Subgraph. In particular, our
 216 results show that, assuming the conjecture on the hardness of 2-CSP that we introduce, for
 217 some constant $0 < \varepsilon \leq 1/2$, there is no efficient $2^{(\log n)^\varepsilon}$ -approximation algorithm for Dense
 218 k -Coloring, unless $\text{NP} \subseteq \text{BPTIME}(n^{O(\log n)})$.

219 Our Conjecture on Hardness of 2-CSP's

220 We now turn to describe our new conjecture on hardness of 2-CSP's. We consider the following
 221 bipartite version of the Constraint Satisfaction Problem with 2 variables per constraint (2-
 222 CSP). The input consists of two sets X and Y of variables, together with an integer $A \geq 1$.
 223 Every variable in $X \cup Y$ takes values in $[A] = \{1, \dots, A\}$. We are also given a collection \mathcal{C}
 224 of constraints, where each constraint $C(x, y) \in \mathcal{C}$ is defined over a pair of variables $x \in X$
 225 and $y \in Y$. For each such constraint, we are given a truth table that, for every pair of
 226 assignments a to x and a' to y , indicates whether (a, a') satisfy the constraint. The value
 227 of the CSP is the largest fraction of constraints that can be simultaneously satisfied by an
 228 assignment to the variables. For given values $0 < s < c \leq 1$, the (c, s) -Gap-CSP problem is

229 the problem of distinguishing CSP's of value at least c from those of value at most s .

230 We can associate, to each constraint $C = C(x, y) \in \mathcal{C}$, a bipartite graph $G_C = (L, R, E)$,
 231 where $L = R = [A]$, and there is an edge (a, a') in E iff the assignments a to x and a' to y
 232 satisfy C . Notice that instance \mathcal{I} of the Bipartite 2-CSP problem is completely defined by
 233 X, Y, A, \mathcal{C} , and the graphs in $\{G_C\}_{C \in \mathcal{C}}$, so we will denote $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$. We
 234 let the *size* of instance \mathcal{I} be $\text{size}(\mathcal{I}) = |\mathcal{C}| \cdot A^2 + |X| + |Y|$. We sometimes refer to A as the
 235 *size of the alphabet for instance \mathcal{I}* . We say that instance \mathcal{I} of 2-CSP is d -to- d' iff for every
 236 constraint C , every vertex of G_C that lies in L has degree at most d , and every vertex that
 237 lies in R has degree at most d' . (We note that this is somewhat different from the standard
 238 definition, that requires that all vertices in L have degree exactly d and all vertices of R have
 239 degree exactly d' . In the standard definition, the alphabet sizes for variables in X and Y
 240 may be different, that is, variables in X take values in $[A]$ and variables of Y take values in
 241 $[A']$ for some integers A, A' . However, this difference is insignificant to our discussion, and it
 242 is more convenient for us to use this slight variation of the standard definition).

243 The famous Unique-Games Conjecture of Khot [35] applies to 1-to-1 CSP's. The conjecture
 244 states that, for any $0 < \varepsilon < 1$, there is a large enough value A , such that the $(1 - \varepsilon, \varepsilon)$ -Gap-CSP
 245 problem is NP-hard for 1-to-1 instances with alphabet size A . The conjecture currently
 246 remains open, though interesting progress has been made on the algorithmic side: the results
 247 of [4] provide an algorithm for the problem with running time $2^{n^{O(1/\varepsilon^{1/3})}}$.

248 A conjecture that is closely related to the Unique-Games Conjecture is the d -to-1 Conjec-
 249 ture of Khot [35]. The conjecture states that, for every $0 < \varepsilon < 1$, and $d > 0$, there is a large
 250 enough value A , such that the $(1, \varepsilon)$ -Gap-CSP problem in d -to-1 instances with alphabet size
 251 A is NP-hard.

252 Håstad [31] proved the following nearly optimal hardness of approximation results for
 253 CSP's: he showed that for every $0 < \varepsilon < 1$, there are values d and A , such that the problem
 254 of $(1, \varepsilon)$ -Gap-CSP in d -to-1 instances with alphabet size A is NP-hard. The value d , however,
 255 depends exponentially on $\text{poly}(1/\varepsilon)$ in this result. In contrast, in the d -to-1 Conjecture, both
 256 d and ε are fixed, and d may not have such a strong dependence on $1/\varepsilon$.

257 On the algorithmic side, the results of [4, 53] provide an algorithm for (c, s) -Gap-CSP
 258 on d -to-1 instances. The running time of the algorithm is $2^{n^{O(1/(\log(1/s))^{1/2})}}$, where the $O(\cdot)$
 259 notation hides factors that are polynomial in d and A .

260 A recent breakthrough in this area is the proof of the 2-to-2 conjecture (now theorem),
 261 that builds on a long sequence of work [5, 6, 22, 23, 37–40]. The theorem proves that for every
 262 $0 < \varepsilon < 1$, there is a large enough value A , such that the $(1 - \varepsilon, \varepsilon)$ -Gap-CSP problem is
 263 NP-hard on 2-to-2 instances with alphabet size A .

264 In this paper, we propose the following conjecture regarding the hardness of Gap-CSP in
 265 d -to- d instances.

266 ► **Conjecture 1.** *There is a constant $0 < \varepsilon \leq 1/2$, such that it is NP-hard to distinguish*
 267 *between $d(n)$ -to- $d(n)$ instances of 2-CSP of size n , that have value at least $1/2$, and those of*
 268 *value at most $s(n)$, where $d(n) = 2^{(\log n)^\varepsilon}$ and $s(n) = 1/2^{64(\log n)^{1/2+\varepsilon}}$.*

269 We now compare this conjecture to existing conjectures and results in this area that we
 270 are aware of. First, in contrast to the d -to-1 conjecture, we allow the parameter d and the
 271 soundness parameter s to be functions of n – the size of the input instance. Note that the
 272 size of the input instance depends on the alphabet size A , so, unlike in the setting of the
 273 d -to-1 conjecture, A may no longer be arbitrarily large compared to d and s .

274 The hardness of approximation result of Håstad [31] for d -to- d CSP's only holds when d
 275 depends exponentially on $\text{poly}(1/s)$, (in particular it may not extend to the setting where

276 $s(n) = 1/2^{64(\log n)^{1/2+\varepsilon}}$, since the size n of the instance depends polynomially on $d(n)$.

277 We can also combine standard constant hardness of approximation results for CSP's (such
278 as, for example, 3-SAT) with the Parallel Repetition theorem, to obtain NP-hardness of
279 $(1, s(n))$ -Gap-CSP on $d(n)$ -to- $d(n)$ instances. Using this approach, if we start from an instance
280 of CSP of size N and a constant hardness gap (with perfect completeness), after ℓ rounds of
281 parallel repetition, we obtain hardness of $(1, s)$ -Gap-CSP on d -to- d instances with $s = 2^{-O(\ell)}$,
282 $d = 2^{O(\ell)}$, and the resulting instance size $n = N^{O(\ell)}$. Note that $d = (1/s)^{\Theta(1)}$ holds,
283 which is different from the relationship between these parameters required by the conjecture.
284 Specifically, by setting the number of repetition to be $\ell = \Theta((\log N)^{(1/2+\varepsilon)/(1/2-\varepsilon)})$, we can
285 ensure the desired bound $s(n) = 1/2^{64(\log n)^{1/2+\varepsilon}}$. However, in this setting, we also get that
286 $d(n) = 2^{\Omega((\log n)^{1/2+\varepsilon})}$, which is significantly higher than the desired value $d(n) = 2^{(\log n)^\varepsilon}$.

287 Lastly, one could attempt to combine the recent proof of the 2-to-2 conjecture with Parallel
288 Repetition in order to reap the benefits of both approaches, but the resulting parameters
289 also fall short of the ones stated in the conjecture.

290 From the above discussion, one can view Conjecture 1 as occupying a middle ground
291 between the d -to-1 conjecture, and the results one can obtain via standard techniques of
292 amplifying a constant hardness of a CSP, such as 3SAT, via Parallel Repetition. We note
293 that, while the conjecture appears closely related to the Unique Games Conjecture and
294 d -to-1 conjecture, we are not aware of any additional formal connections, except for those
295 mentioned above.

296 We now proceed to discuss our results and techniques in more detail.

297 1.1 A More Detailed Overview of our Results and Techniques

298 In addition to posing Conjecture 1 that we already described above, we prove conditional
299 hardness of approximation of the four problems that we consider. We also prove that all
300 four problems are roughly equivalent approximation-wise. We now discuss the conditional
301 hardness of approximation for Densest k -Subgraph and the connections between the four
302 problems that we establish.

303 Conditional Hardness of Densest k -Subgraph.

304 Our first result is a conditional hardness of Densest k -Subgraph. Specifically, we prove that,
305 assuming that Conjecture 1 holds and that $P \neq NP$, for some $0 < \varepsilon \leq 1/2$, there is no efficient
306 approximation algorithm for Densest k -Subgraph problem that achieves approximation factor
307 $2^{(\log N)^\varepsilon}$, where N is the number of vertices in the input graph.

308 We now provide a brief overview of our techniques. The proof of the above result employs
309 a Cook-type reduction, and follows some of the ideas that were introduced in [18]. We
310 assume for contradiction that there is a factor- α algorithm \mathcal{A} for the Densest k -Subgraph
311 problem, where $\alpha = 2^{(\log N)^\varepsilon}$. Given an input instance \mathcal{I} of the 2-CSP problem of size n , that
312 is a $d(n)$ -to- $d(n)$ instance, we construct a constraint graph H representing \mathcal{I} . We gradually
313 decompose graph H into a collection \mathcal{H} of disjoint subgraphs, such that, for each subgraph
314 $H' \in \mathcal{H}$, we can either certify that the value of the corresponding instance of 2-CSP is at
315 most $1/4$, or it is at least β , for some carefully chosen parameter β . In order to compute
316 the decomposition, we start with $\mathcal{H} = \{H\}$. If, for a graph $H' \in \mathcal{H}$, we certified that the
317 corresponding instance of 2-CSP has value at most $1/4$, or at least β , then we say that graph
318 H' is *inactive*. Otherwise, we say that it is *active*. As long as \mathcal{H} contains at least one active
319 graph, we perform iterations. In each iteration, we select an arbitrary active graph $H' \in \mathcal{H}$
320 to process. In order to process H' , we consider an *assignment graph* G' associated with H' ,

321 that contains a vertex for every variable-assignment pair (x, a) , where x is a variable whose
 322 corresponding vertex belongs to H' . We view G' as an instance of the Densest k -Subgraph
 323 problem, for an appropriately chosen parameter k , and apply the approximation algorithm \mathcal{A}
 324 for Densest k -Subgraph to it. Let S be the set of vertices of G' that Algorithm \mathcal{A} computes
 325 as a solution to this instance. Note that S is a set of vertices in the assignment graph G' ,
 326 while \mathcal{H} is a family of subgraphs of the constraint graph H . We exploit the set S of vertices
 327 in order to either (i) compute a large subset $E' \subseteq E(H')$ of edges, such that, if we denote by
 328 $\mathcal{C}' \subseteq \mathcal{C}$ the set of constraints corresponding to E' , then at most $1/4$ of the constraints of \mathcal{C}'
 329 can be simultaneously satisfied; or (ii) compute a large subset $E' \subseteq E(H')$ of edges as above,
 330 and certify that at least a β -fraction of such constraints can be satisfied; or (iii) compute a
 331 subgraph $H'' \subseteq H'$, such that $|V(H'')| \ll |V(H')|$, and the number of edges contained in
 332 graphs H'' and $H' \setminus V(H'')$ is sufficiently large compared to $E(H')$. In the former two cases,
 333 we replace H' with graph $H'[E']$ in \mathcal{H} , and graph $H'[E']$ becomes inactive. In the latter case,
 334 we replace H' with two graphs: H'' and $H' \setminus V(H'')$, that both remain active. The algorithm
 335 terminates once every graph in \mathcal{H} is inactive. The crux of the analysis of the algorithm is to
 336 show that, when the algorithm terminates, the total number of edges lying in the subgraphs
 337 $H' \in \mathcal{H}$ is high, compared to $|E(H)|$. The specific fraction of edges that remain in the
 338 subgraphs $H' \in \mathcal{H}$ is governed by the parameters $s(n)$ and $d(n)$, and the specific relationship
 339 between these parameters in Conjecture 1 is selected to ensure that many edges remain in the
 340 graphs of \mathcal{H} when the algorithm terminates. The algorithm for decomposing graph H into
 341 subgraphs and its analysis employ some of the techniques and ideas introduced in [18], and
 342 is very similar in spirit to the hardness of approximation proof of the (r,h)-Graph Partitioning
 343 with Bundles problem, though details are different. We employ this decomposition algorithm
 344 multiple times, in order to obtain a partition (E_0, E_1, \dots, E_z) of the set $E(H)$ of edges
 345 of the constraint graph into a small number of subsets, such that, among the constraints
 346 corresponding to the edges of E_0 , at most a $1/4$ -fraction can be satisfied by any assignment
 347 to $X \cup Y$, and, for all $1 \leq i \leq z$, a large fraction of constraints corresponding to edges of E_i
 348 can be satisfied by some assignment. Depending on the cardinality of the set E_0 of edges we
 349 then determine whether \mathcal{I} is a YES-INSTANCE or a NO-INSTANCE.

350 Reductions from Dense k -Coloring and (r,h)-Graph Partitioning to Densest k -Subgraph.

351 We show that, if there is an efficient factor $\alpha(n)$ -approximation algorithm for the Densest
 352 k -Subgraph problem, then there is a randomized efficient $O(\alpha(n^2) \cdot \text{poly log } n)$ -approximation
 353 algorithm for Dense k -Coloring, and a randomized efficient $O(\alpha(n^2) \cdot \text{poly log } n)$ -approximation
 354 algorithm for (r,h)-Graph Partitioning. The two reductions are very similar, so we focus on
 355 describing the first one. We believe that the reduction is of independent interest, and uses
 356 unusual techniques.

357 We assume that there is an $\alpha(n)$ -approximation algorithm for the Densest k -Subgraph
 358 problem. In order to obtain an approximation algorithm for Dense k -Coloring, we start by
 359 formulating a natural LP-relaxation for the problem. Unfortunately, this LP-relaxation has
 360 a large number of variables: roughly $n^{\Theta(k)}$, where n is the number of vertices in the input
 361 graph and k is the parameter of the Dense k -Coloring problem instance. We then show an
 362 efficient algorithm, that, given a solution to the LP-relaxation, whose support size is bounded
 363 by $\text{poly}(n)$, computes an approximate integral solution to the Dense k -Coloring problem.

364 The main challenge is that, since the LP relaxation has $n^{\Theta(k)}$ variables, it is unclear
 365 how to solve it efficiently. We consider the dual linear program, that has $\text{poly}(n)$ variables
 366 and $n^{\Theta(k)}$ constraints. Using the $\alpha(n)$ -approximation algorithm for Densest k -Subgraph as
 367 a subroutine, we design an approximate separation oracle for the dual LP, that allows us

368 to solve the original LP-relaxation for Dense k -Coloring, obtaining a solution whose support
 369 size is bounded by $\text{poly}(n)$. By applying the LP-rounding approximation algorithm to this
 370 solution, we obtain the desired approximate solution to the input instance of Dense k -Coloring.

371 **Reductions from Densest k -Subgraph to (r,h)-Graph Partitioning and Dense k -Coloring.**

372 We prove that, if there is an efficient $\alpha(n)$ -approximation algorithm for Dense k -Coloring,
 373 then there is a randomized algorithm for the Densest k -Subgraph problem, whose running
 374 time is $n^{O(\log n)}$, that with high probability obtains an $O(\alpha(n^{O(\log n)}) \cdot \log n)$ -approximate
 375 solution to the input instance of the problem. We also show a similar reduction from Densest
 376 k -Subgraph to (r,h)-Graph Partitioning, but now the resulting approximation factor for Densest
 377 k -Subgraph becomes $O((\alpha(n^{O(\log n)}))^3 \cdot \log^2 n)$. By combining these reductions with our
 378 conditional hardness result for Densest k -Subgraph, we get that, assuming Conjecture 1,
 379 for some constant $0 < \varepsilon \leq 1/2$, there is no efficient $2^{(\log n)^\varepsilon}$ -approximation algorithm for
 380 (r,h)-Graph Partitioning and for Dense k -Coloring, unless $\text{NP} \subseteq \text{BPTIME}(n^{O(\log n)})$.

381 The two reductions are very similar; we focus on the reduction to Dense k -Coloring in
 382 this overview. Our construction is inspired by the results of [34], and we borrow some of our
 383 ideas from them. Assume that there is an efficient $\alpha(n)$ -approximation algorithm for Dense
 384 k -Coloring. Let G be an instance of the Densest k -Subgraph problem. The main difficulty in
 385 the reduction is that it is possible that G only contains one very dense subgraph induced by
 386 k vertices, while the Dense k -Coloring problem requires that the input graph G can essentially
 387 be partitioned into many such dense subgraphs. To overcome this difficulty, we construct
 388 a random “inflated” bipartite graph H , that contains $n^{O(\log n)}$ vertices, where $n = |V(G)|$.
 389 Every vertex of G is mapped to some vertex of H at random, while every edge of G is
 390 mapped to a large number of edges of H . This allows us to ensure that, if G contains a
 391 subgraph G' induced by a set of k vertices, where $|E(G')| = R$, then graph H can essentially
 392 be partitioned into a large number of subgraphs that contain k vertices each, and many of
 393 them contain close to R edges. Therefore, we can apply our $\alpha(n)$ -approximation algorithm
 394 for Dense k -Coloring to the new graph H . The main challenge in the reduction is that,
 395 while this approximation algorithm is guaranteed to return a large number of disjoint dense
 396 subgraphs of H , since every edge of G contributes many copies to H , it is not clear that one
 397 can extract a single dense subgraph of G from dense subgraphs of H . The main difficulty in
 398 the reduction is to ensure that, on the one hand, a single k -vertex dense subgraph in G can
 399 be translated into $|V(H)|/k$ dense subgraphs of H ; and, on the other hand, a dense k -vertex
 400 subgraph of H can be translated into a dense subgraph of G on k vertices. We build on and
 401 expand the ideas from [34] in order to ensure these properties.

402 **Reductions between (r,h)-Graph Partitioning and Maximum Bounded-Crossing Subgraph.**

403 Lastly, we provide reductions between (r,h)-Graph Partitioning and Maximum Bounded-Crossing
 404 Subgraph in both directions. First, we show that, if there is an efficient factor $\alpha(n)$ -
 405 approximation algorithm for (r,h)-Graph Partitioning, then there is an efficient $O(\alpha(n) \cdot$
 406 $\text{poly} \log n)$ -approximation algorithm for Maximum Bounded-Crossing Subgraph. On the other
 407 hand, an efficient $\alpha(n)$ -approximation algorithm for Maximum Bounded-Crossing Subgraph
 408 implies an efficient $O((\alpha(n))^2 \cdot \text{poly} \log n)$ -approximation algorithm for (r,h)-Graph Partitioning.
 409 Combined with our conditional hardness of approximation for (r,h)-Graph Partitioning, we
 410 get that, assuming Conjecture 1, for some constant $0 < \varepsilon \leq 1/2$, there is no efficient
 411 $2^{(\log n)^\varepsilon}$ -approximation algorithm for Maximum Bounded-Crossing Subgraph, unless $\text{NP} \subseteq$
 412 $\text{BPTIME}(n^{O(\log n)})$.

Both these reductions exploit the following connection between crossing number and graph partitioning: if a graph G has a drawing with at most L crossings, then there is a balanced cut in G , containing at most $O\left(\sqrt{L + \Delta \cdot |E(G)|}\right)$ edges, where Δ is maximum vertex degree in G . This result can be viewed as an extension of the classical Planar Separator Theorem of [45]. Another useful fact exploited in both reductions is that any graph G with m edges has a plane drawing with at most m^2 crossings. In particular, if $\mathcal{H} = \{H_1, \dots, H_r\}$ is a solution to an instance of the (r, h) -Graph Partitioning problem on graph G , then there is a drawing of graph $H = \bigcup_{i=1}^r H_i$, in which the number of crossings is bounded by $r \cdot h^2$. These two facts establish a close relationship between the (r, h) -Graph Partitioning and Maximum Bounded-Crossing Subgraph problems, that are exploited in both our reductions.

We have now obtained a chain of reductions, showing that all four problems, Densest k -Subgraph, Dense k -Coloring, (r, h) -Graph Partitioning, and Maximum Bounded-Crossing Subgraph are almost equivalent from approximation viewpoint, if we consider sufficiently large approximation factors and allow randomized quasi-polynomial time algorithms. We also obtain conditional hardness of approximation results for all four problems based on Conjecture 1.

Organization.

We start with preliminaries in Section 2. In Section 3 we provide the conditional hardness of approximation proof for the Densest k -Subgraph problem. In Section 4 we provide our reductions from Dense k -Coloring and (r, h) -Graph Partitioning to Densest k -Subgraph, and in Section 5 we provide reductions in the opposite direction. Lastly, in Section 6 we provide reductions between (r, h) -Graph Partitioning and Maximum Bounded-Crossing Subgraph. Due to lack of space, some of the proofs are deferred to the full version of the paper.

2 Preliminaries

By default, all logarithms are to the base of 2. For a positive integer N , we denote by $[N] = \{1, 2, \dots, N\}$. All graphs are finite, simple and undirected. We say that an event holds with high probability if the probability of the event is $1 - 1/n^c$ for a large enough constant c , where n is the number of vertices in the input graph.

2.1 General Notation

Let G be a graph and let S be a subset of its vertices. We denote by $G[S]$ the subgraph of G induced by S . For two disjoint subsets A, B of vertices of G , we denote by $E_G(A, B)$ the set of all edges with one endpoint in A and the other endpoint in B , and we denote by $E_G(A)$ the set of all edges with both endpoints in A . Given a graph G and a vertex $v \in V(G)$, we denote by $\deg_G(v)$ the degree of v in G . For a subset S of vertices of G , its *volume* is $\text{vol}_G(S) = \sum_{v \in S} \deg_G(v)$. We sometimes omit the subscript G if it is clear from the context.

Given a graph G , a drawing φ of G is an embedding of G into the plane, that maps every vertex v of G to a point (called the *image of v* and denoted by $\varphi(v)$), and every edge e of G to a simple curve (called the *image of e* and denoted by $\varphi(e)$), that connects the images of its endpoints. If e is an edge of G and v is a vertex of G , then the image of e may only contain the image of v if v is an endpoint of e . Furthermore, if some point p belongs to the images of three or more edges of G , then p must be the image of a common endpoint of all edges e with $p \in \varphi(e)$. We say that two edges e, e' of G *cross* at a point p , if $p \in \varphi(e) \cap \varphi(e')$, and p is not the image of a shared endpoint of these edges. Given a graph G and a drawing

456 φ of G in the plane, we use $\text{cr}(\varphi)$ to denote the number of crossings in φ , and the crossing
 457 number of G , denoted by $\text{CrN}(G)$, is the minimum number of crossings in any drawing of G .

458 2.2 Problem Definitions and Additional Notation

459 In this paper we consider the following four problems: Densest k -Subgraph, Dense k -Coloring,
 460 (r, h) -Graph Partitioning and Maximum Bounded-Crossing Subgraph. We now define the
 461 problems, along with some additional notation.

462 Densest k -Subgraph.

463 In the Densest k -Subgraph problem, the input is a graph G and an integer $k > 0$. The goal is
 464 to compute a subset $S \subseteq V(G)$ of k vertices, maximizing $|E_G(S)|$. We denote an instance
 465 of the problem by $\text{DkS}(G, k)$, and we denote the value of the optimal solution to instance
 466 $\text{DkS}(G, k)$ by $\text{OPT}_{\text{DkS}}(G, k)$.

467 We also consider a bipartite version of the Densest k -Subgraph problem, called
 468 Bipartite Densest (k_1, k_2) -Subgraph. This problem was first studied in [2]. The input to the
 469 problem is a bipartite graph $G = (A, B, E)$ and positive integers k_1, k_2 . The goal is to
 470 compute a subset $S \subseteq V(G)$ of vertices with $|S \cap A| = k_1$ and $|S \cap B| = k_2$, such that $|E_G(S)|$
 471 is maximized. An instance of this problem is denoted by $\text{BDkS}(G, k_1, k_2)$, and the value of
 472 the optimal solution to instance $\text{BDkS}(G, k_1, k_2)$ is denoted by $\text{OPT}_{\text{BDkS}}(G, k_1, k_2)$. The
 473 following lemma shows that the Bipartite Densest (k_1, k_2) -Subgraph problem and the Densest
 474 k -Subgraph problem are roughly equivalent from the approximation viewpoint. Similar results
 475 were also shown in prior work.

476 ► **Lemma 2.** *Let $\alpha : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be an increasing function such that $\alpha(n) = o(n)$. Then the*
 477 *following hold:*

478 ■ *If there exists an $\alpha(n)$ -approximation algorithm for the Densest k -Subgraph problem with*
 479 *running time at most $T(n)$, where n is the number of vertices in the input graph, then there*
 480 *exists an $O(\alpha(N^2))$ -approximation algorithm for the Bipartite Densest (k_1, k_2) -Subgraph*
 481 *problem, with running time $O(T(N^2) \cdot \text{poly}(N))$, where N is the number of vertices in*
 482 *the input graph. Moreover, if the algorithm for Densest k -Subgraph is deterministic, then*
 483 *so is the algorithm for Bipartite Densest (k_1, k_2) -Subgraph.*

484 ■ *Similarly, if there exists an efficient $\alpha(N)$ -approximation algorithm for the Bipartite*
 485 *Densest (k_1, k_2) -Subgraph problem, where N is the number of vertices in the input graph,*
 486 *then there exists an efficient $O(\alpha(2n))$ -approximation algorithm for the Densest k -Subgraph*
 487 *problem, where n is the number of vertices in the input graph. Moreover, if the algorithm*
 488 *for Bipartite Densest (k_1, k_2) -Subgraph is deterministic, then so is the algorithm for*
 489 *Densest k -Subgraph.*

490 Dense k -Coloring.

491 The input to the Dense k -Coloring problem consists of an n -vertex graph G and an integer
 492 $k > 0$, such that n is an integral multiple of k . The goal is to compute a partition of $V(G)$
 493 into n/k subsets $S_1, \dots, S_{n/k}$ of cardinality k each, while maximizing $\sum_{i=1}^{n/k} |E_G(S_i)|$. An
 494 instance of the Dense k -Coloring problem is denoted by $\text{DkC}(G, k)$, and the value of the
 495 optimal solution to instance $\text{DkC}(G, k)$ is denoted by $\text{OPT}_{\text{DkC}}(G, k)$.

17:12 A New Conjecture on 2-CSP's with Implications to Densest k -Subgraph

496 (r, h) -Graph Partitioning.

497 The input to the (r, h) -Graph Partitioning problem consists of a graph G , and integers $r, h > 0$.
498 The goal is to compute r vertex-disjoint subgraphs H_1, \dots, H_r of G , such that for all
499 $1 \leq i \leq r$, $|E(H_i)| \leq h$, while maximizing $\sum_{i=1}^r |E(H_i)|$. An instance of the (r, h) -Graph
500 Partitioning problem is denoted by $\text{GP}(G, r, h)$, and the value of the optimal solution to
501 instance $\text{GP}(G, r, h)$ is denoted by $\text{OPT}_{\text{GP}}(G, r, h)$.

502 Maximum Bounded-Crossing Subgraph.

503 In the Maximum Bounded-Crossing Subgraph problem, the input is a graph G and an integer
504 $L > 0$. The goal is to compute a subgraph $H \subseteq G$ with $\text{CrN}(H) \leq L$, while maximizing
505 $|E(H)|$. An instance of the Maximum Bounded-Crossing Subgraph problem is denoted by
506 $\text{MBCS}(G, L)$, and the value of the optimal solution to instance $\text{MBCS}(G, L)$ is denoted by
507 $\text{OPT}_{\text{MBCS}}(G, L)$. We note that we can assume that $L \leq |V(G)|^4$, as otherwise the optimal
508 solution is the whole graph G , since the crossing number of a simple graph G is at most
509 $|E(G)|^2 \leq |V(G)|^4$.

510 **3** Conditional Hardness of Densest k -Subgraph

511 3.1 Conjecture on Hardness of 2-CSP's

512 We consider the Bipartite 2-CSP problem, that is defined as follows. The input to the problem
513 consists of two sets X, Y of variables, together with an integer $A > 1$. Every variable
514 $z \in X \cup Y$ takes values in set $[A] = \{1, \dots, A\}$. We are also given a collection \mathcal{C} of constraints,
515 where each constraint $C(x, y) \in \mathcal{C}$ is defined over a pair of variables $x \in X$ and $y \in Y$. For
516 each such constraint, we are given a truth table that, for every pair of assignments a to x
517 and a' to y , specifies whether (a, a') satisfy constraint $C(x, y)$. The value of the CSP is the
518 largest fraction of constraints that can be simultaneously satisfied by an assignment to the
519 variables.

520 We associate with each constraint $C = C(x, y) \in \mathcal{C}$, a bipartite graph $G_C = (L, R, E)$,
521 where $L = R = [A]$, and there is an edge (a, a') in E iff the assignments a to x and a' to y
522 satisfy C . Notice that instance \mathcal{I} of the Bipartite 2-CSP problem is completely determined by
523 X, Y, A, \mathcal{C} , and the graphs in $\{G_C\}_{C \in \mathcal{C}}$, so we will denote $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$. The
524 size of instance \mathcal{I} is defined to be $\text{size}(\mathcal{I}) = |\mathcal{C}| \cdot A^2 + |X| + |Y|$.

525 Consider some instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP. We say that \mathcal{I} is a
526 d -to- d instance if, for every constraint C , every vertex of graph $G_C = (L, R, E)$ has degree
527 at most d .

528 Consider now some functions $d(n), s(n) : \mathbb{R}^+ \rightarrow \mathbb{R}^+$. We assume that, for all n ,
529 $d(n) \geq 1$ and $s(n) < 1$. In a $(d(n), s(n))$ -2CSP problem, the input is an instance $\mathcal{I} =$
530 $(X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP, such that, if we denote by $n = \text{size}(\mathcal{I})$, then the
531 instance is $d(n)$ -to- $d(n)$. We say that \mathcal{I} is a YES-INSTANCE, if there is some assignment
532 to the variables of $X \cup Y$ that satisfies at least $|\mathcal{C}|/2$ of the constraints, and we say that
533 it is a NO-INSTANCE, if the largest number of constraints of \mathcal{C} that can be simultaneously
534 satisfied by any assignment is at most $s(n) \cdot |\mathcal{C}|$. Given an instance \mathcal{I} of $(d(n), s(n))$ -2CSP
535 problem, the goal is to distinguish between the case where \mathcal{I} is a YES-INSTANCE and the
536 case where \mathcal{I} is a NO-INSTANCE. If \mathcal{I} is neither a YES-INSTANCE nor a NO-INSTANCE, the
537 output of the algorithm can be arbitrary. We now state our conjecture regarding hardness of
538 $(d(n), s(n))$ -2CSP, that is a restatement of Conjecture 1 from the Introduction.

539 ► **Conjecture 3.** *There is a constant $0 < \varepsilon \leq 1/2$, such that the $(d(n), s(n))$ -2CSP problem*
 540 *is NP-hard for $d(n) = 2^{(\log n)^\varepsilon}$ and $s(n) = 1/2^{64(\log n)^{1/2+\varepsilon}}$.*

541 3.2 Conditional Hardness of Densest k -Subgraph

542 In the remainder of this section, we prove the following theorem on the conditional hardness
 543 of Densest k -Subgraph.

544 ► **Theorem 4.** *Assume that Conjecture 3 holds and that $P \neq NP$. Then for some $0 < \varepsilon \leq 1/2$,*
 545 *there is no efficient approximation algorithm for Densest k -Subgraph problem that achieves*
 546 *approximation factor $2^{(\log N)^\varepsilon}$, where N is the number of vertices in the input graph.*

547 In fact we will prove a slightly more general theorem, that will be useful for us later.

548 ► **Theorem 5.** *Suppose there is an algorithm for the Densest k -Subgraph problem, that,*
 549 *given an instance $\text{DkS}(G, k)$ with $|V(G)| = N$, in time at most $T(N)$, computes a factor*
 550 *$2^{(\log N)^\varepsilon}$ -approximate solution to the problem, for some constant $0 < \varepsilon \leq 1/2$. Then there*
 551 *is an algorithm, that, given an instance \mathcal{I} of $(d(n), s(n))$ -2CSP problem of size n , where*
 552 *$d(n) = 2^{(\log n)^\varepsilon}$ and $s(n) = 1/2^{64(\log n)^{1/2+\varepsilon}}$, responds “YES” or “NO”, in time $O(\text{poly}(n) \cdot$
 553 $T(\text{poly}(n)))$. If \mathcal{I} is a YES-INSTANCE, the algorithm is guaranteed to respond “YES”, and if*
 554 *it is a NO-INSTANCE, it is guaranteed to respond “NO”.*

555 Theorem 4 immediately follows from Theorem 5. The remainder of this section is
 556 dedicated to proving Theorem 5. A central notion that we use is a *constraint graph* that is
 557 associated with an instance \mathcal{I} of 2-CSP.

558 Let $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ be an instance of the Bipartite 2-CSP problem. The
 559 *constraint graph* associated with instance \mathcal{I} is denoted by $H(\mathcal{I})$, and it is defined as follows.
 560 The set of vertices of $H(\mathcal{I})$ is the union of two subsets: set $V = \{v(x) \mid x \in X\}$ of vertices
 561 representing the variables of X , and set $U = \{v(y) \mid y \in Y\}$ of vertices representing the
 562 variables of Y . For convenience, we will not distinguish between the vertices of V and the
 563 variables of X , so we will identify each variable $x \in X$ with its corresponding vertex $v(x)$.
 564 Similarly, we will not distinguish between vertices of U and variables of Y . The set of
 565 edges of $H(\mathcal{I})$ contains, for every constraint $C = C(x, y) \in \mathcal{C}$, edge $e_C = (x, y)$. We say
 566 that edge e_C *represents* the constraint C . Notice that, if E' is a subset of edges of $H(\mathcal{I})$,
 567 then we can define a set $\Phi(E') \subseteq \mathcal{C}$ of constraints that the edges of E' represent, namely:
 568 $\Phi(E') = \{C \in \mathcal{C} \mid e_C \in E'\}$. Next, we define bad sets of constraints and bad sets of edges.

569 ► **Definition 6** (Bad Set of Constraints and Bad Collection of Edges). *Let $\mathcal{C}' \subseteq \mathcal{C}$ be a collection*
 570 *of constraints of an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP. We say that \mathcal{C}' is*
 571 *a bad set of constraints if the largest number of constraints of \mathcal{C}' that can be simultaneously*
 572 *satisfied by any assignment to the variables of $X \cup Y$ is at most $\frac{|\mathcal{C}'|}{4}$. If $E' \subseteq E(H(\mathcal{I}))$ is a*
 573 *set of edges of $H(\mathcal{I})$, whose corresponding set $\Phi(E')$ of constraints is bad, then we say that*
 574 *E' is a bad collection of edges.*

575 The next observation easily follows from the definition of a bad set of constraints.

576 ► **Observation 7.** *Let $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ be an instance of bipartite 2-CSP, and let*
 577 *$\mathcal{C}', \mathcal{C}'' \subseteq \mathcal{C}$ be two disjoint sets of constraints that are both bad. Then $\mathcal{C}' \cup \mathcal{C}''$ is also a bad set*
 578 *of constraints.*

579 Next, we define good subsets of constraints and good subgraphs of the constraint graph
 580 $H(\mathcal{I})$.

17:14 A New Conjecture on 2-CSP's with Implications to Densest k -Subgraph

581 ► **Definition 8** (Good Set of Constraints and Good Subgraphs of $H(\mathcal{I})$). Let $\mathcal{C}' \subseteq \mathcal{C}$ be a
582 collection of constraints of an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP, and
583 let $0 < \beta \leq 1$ be a parameter. We say that \mathcal{C}' is a β -good set of constraints, if there is an
584 assignment to variables of $X \cup Y$ that satisfies at least $\frac{|\mathcal{C}'|}{\beta}$ constraints of \mathcal{C}' . If $E' \subseteq E(H(\mathcal{I}))$
585 is a set of edges of $H(\mathcal{I})$, whose corresponding set $\Phi(E')$ of constraints is β -good, then we say
586 that E' is a β -good collection of edges. Lastly, if $H' \subseteq H(\mathcal{I})$ is a subgraph of the constraint
587 graph, and the set $E(H')$ of edges is β -good, then we say that graph H' is β -good.

588 The next observation easily follows from the definition of a good set of constraints.

589 ► **Observation 9.** Let $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ be an instance of bipartite 2-CSP, let
590 $0 < \beta \leq 1$ be a parameter, and let H', H'' be two subgraphs of $H(\mathcal{I})$ that are both β -good and
591 disjoint in their vertices. Then graph $H' \cup H''$ is also β -good.

592 The observation follows from the fact that, since graphs H', H'' are disjoint in their
593 vertices, if we let $\mathcal{C}' = \Phi(E(H'))$, $\mathcal{C}'' = \Phi(E(H''))$ be the sets of constraints associated with
594 the edge sets of both graphs, then the variables participating in the constraints of \mathcal{C}' are
595 disjoint from the variables participating in the constraints of \mathcal{C}'' .

596 The following theorem is key in proving Theorem 5.

597 ► **Theorem 10.** Assume that there exists a constant $0 < \varepsilon \leq 1/2$, and an $\alpha(N)$ -approximation
598 algorithm \mathcal{A} for the Densest k -Subgraph problem, whose running time is at most $T(N)$, where
599 N is the number of vertices in the input graph, and $\alpha(N) = 2^{(\log N)^\varepsilon}$. Then there is an
600 algorithm, whose input consists of an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP
601 and parameter n that is greater than a large enough constant, so that $\text{size}(\mathcal{I}) \leq n$ holds, and
602 \mathcal{I} is a $d(n)$ -to- $d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^\varepsilon}$. Let $\beta = 2^{8(\log n)^{1/2+\varepsilon}}$,
603 and let $r = \lceil \beta \cdot \log n \rceil$. The algorithm returns a partition (E^b, E_1, \dots, E_r) of $E(H(\mathcal{I}))$, such
604 that E^b is a bad set of edges, and for all $1 \leq i \leq r$, set E_i of edges is β^3 -good. The running
605 time of the algorithm is $O(T(\text{poly}(n)) \cdot \text{poly}(n))$.

606 The proof of Theorem 5 easily follows from Theorem 10. Assume that there exists a
607 constant $0 < \varepsilon \leq 1/2$, and an $\alpha(N)$ -approximation algorithm \mathcal{A} for the Densest k -Subgraph
608 problem, whose running time is at most $T(N)$, where N is the number of vertices in the
609 input graph, and $\alpha(N) = 2^{(\log N)^\varepsilon}$. We show an algorithm for the $(d(n), s(n))$ -2CSP problem,
610 for $d(n) = 2^{(\log n)^\varepsilon}$ and $s(n) = 1/2^{64(\log n)^{1/2+\varepsilon}}$. Let $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ be an input
611 instance of the Bipartite 2-CSP problem, with $\text{size}(\mathcal{I}) = n$, so that \mathcal{I} is a $d(n)$ -to- $d(n)$ instance
612 of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^\varepsilon}$. If n is bounded by a constant, then we can determine
613 whether \mathcal{I} is a YES-INSTANCE or a NO-INSTANCE by exhaustively trying all assignments to
614 its variables. Therefore, we assume that n is greater than a large enough constant. We apply
615 the algorithm from Theorem 10 to this instance \mathcal{I} . Let (E^b, E_1, \dots, E_r) be the partition of
616 the edges of $E(H(\mathcal{I}))$ that the algorithm returns. We now consider two cases.

617 Assume first that $|E^b| > 2|\mathcal{C}|/3$. Let $\mathcal{C}^b \subseteq \mathcal{C}$ be the set of all constraints that correspond
618 to the edges of E^b . Recall that set \mathcal{C}^b of constraints is bad, so in any assignment, at most
619 $\frac{|\mathcal{C}^b|}{4}$ of the constraints in \mathcal{C}^b may be satisfied. Therefore, if f is any assignment to variables
620 of $X \cup Y$, the number of constraints in \mathcal{C} that are not satisfied by f is at least $\frac{3|\mathcal{C}^b|}{4} > \frac{|\mathcal{C}|}{2}$.
621 Clearly, \mathcal{I} may not be a YES-INSTANCE in this case. Therefore, if $|E^b| > 2|\mathcal{C}|/3$, we report
622 that \mathcal{I} is a NO-INSTANCE.

623 If $|E^b| \leq 2|\mathcal{C}|/3$, then we report that \mathcal{I} is a YES-INSTANCE. It is now enough to show
624 that, if $|E^b| \leq 2|\mathcal{C}|/3$, then instance \mathcal{I} may not be a NO-INSTANCE. In other words, it is
625 enough to show that there is an assignment that satisfies more than $\frac{|\mathcal{C}|}{2^{64(\log n)^{1/2+\varepsilon}}}$ constraints.

626 Indeed, since $|E^b| \leq 2|\mathcal{C}|/3$, there is an index $1 \leq i \leq r$, with $|E_i| \geq \frac{|\mathcal{C}|}{3r}$. Since set E_i of
 627 edges is β^3 -good, there is an assignment to the variables of $X \cup Y$, that satisfies at least
 628 $\frac{|E_i|}{\beta^3} \geq \frac{|\mathcal{C}|}{3r\beta^3}$ constraints that correspond to the edges of E_i . Recall that $\beta = 2^{8(\log n)^{1/2+\epsilon}}$
 629 and $r = \lceil \beta \cdot \log n \rceil$. Therefore, $3r\beta^3 \leq 6\beta^4 \log n \leq 2^{64(\log n)^{1/2+\epsilon}}$. We conclude that there is
 630 an assignment satisfying at least $|\mathcal{C}|/2^{64(\log n)^{1/2+\epsilon}}$ constraints, and so \mathcal{I} may not be a NO-
 631 INSTANCE. It is easy to verify that the running time of the algorithm is $O(T(\text{poly}(n)) \cdot \text{poly}(n))$.

632 To conclude, we have shown that, if there is an $\alpha(N)$ -approximation algorithm \mathcal{A} for
 633 the Densest k -Subgraph problem, with running time at most $T(N)$, where N is the number
 634 of vertices in the input graph, and $\alpha(N) = 2^{(\log N)^\epsilon}$, then there is an algorithm for the
 635 $(d(n), s(n))$ -2CSP problem, for $d(n) = 2^{(\log n)^\epsilon}$ and $s(n) = 1/2^{8(\log n)^{1/2+\epsilon}}$, whose running
 636 time is $O(T(\text{poly}(n)) \cdot \text{poly}(n))$.

637 In the remainder of this section we prove Theorem 10.

638 3.3 Proof of Theorem 10

639 The following theorem is the main technical ingredient of the proof of Theorem 10.

640 ► **Theorem 11.** *Assume that there exists an $\alpha(N)$ -approximation algorithm \mathcal{A} for the*
 641 *Bipartite Densest (k_1, k_2) -Subgraph problem, whose running time is at most $T(N)$, where*
 642 *N is the number of vertices in the input graph. Then there is an algorithm, that, given*
 643 *an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP and parameters $n, \beta \geq 1$, so that*
 644 *size $(\mathcal{I}) \leq n$, $\beta \geq 2^{30}(\alpha(n))^3(\log n)^{12}$, and \mathcal{I} is a $d(n)$ -to- $d(n)$ instance of Bipartite 2-CSP,*
 645 *for some function $d(n)$, in time $O(T(n) \cdot \text{poly}(n))$, does one of the following:*

- 646 ■ *either correctly establishes that graph $H(\mathcal{I})$ is β^3 -good; or*
- 647 ■ *computes a bad set $\mathcal{C}' \subseteq \mathcal{C}$ of constraints, with $|\mathcal{C}'| \geq \frac{|\mathcal{C}|}{8 \log^2 n}$; or*
- 648 ■ *computes a subgraph $H' = (X', Y', E')$ of $H(\mathcal{I})$, for which the following hold:*
 - 649 ■ $|X'| \leq \frac{2d(n) \cdot |X|}{\beta}$;
 - 650 ■ $|Y'| \leq \frac{2d(n) \cdot |Y|}{\beta}$; and
 - 651 ■ $|E'| \geq \frac{\text{vol}_H(X' \cup Y')}{2048d(n) \cdot \alpha(n) \cdot \log^4 n}$.

652 The proof of Theorem 11 partially relies on ideas and techniques from [18], and is deferred
 653 to the full version of the paper. We now complete the proof of Theorem 10 using Theorem 11,
 654 starting with the following simple corollary, whose proof is deferred to the full version of the
 655 paper.

656 ► **Corollary 12.** *Assume that there exists an $\alpha(N)$ -approximation algorithm \mathcal{A} for the Bipartite*
 657 *Densest (k_1, k_2) -Subgraph problem, whose running time is at most $T(N)$, where N is the*
 658 *number of vertices in the input graph. Then there is an algorithm, whose input consists of*
 659 *an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP and parameters $n, \beta \geq 1$, so that*
 660 *size $(\mathcal{I}) \leq n$, $\beta \geq 2^{30}(\alpha(n))^3(\log n)^{12}$, and \mathcal{I} is a $d(n)$ -to- $d(n)$ instance of Bipartite 2-CSP.*
 661 *The algorithm returns a partition (E_1, E_2) of $E(H(\mathcal{I}))$, where E_1 is a bad set of edges, and:*

- 662 ■ *either the algorithm correctly certifies that E_2 is a β^3 -good set of edges; or*
- 663 ■ *it computes a subgraph $H' = (X', Y', E')$ of $H(\mathcal{I})$, with $E(H') \subseteq E_2$, for which the*
 664 *following hold:*
 - 665 ■ $|X'| \leq \frac{2d(n) \cdot |X|}{\beta}$;
 - 666 ■ $|Y'| \leq \frac{2d(n) \cdot |Y|}{\beta}$; and
 - 667 ■ $|E'| \geq \frac{|E_2^*|}{2048d(n) \cdot \alpha(n) \cdot \log^4 n}$, where E_2^* is a set of edges containing every edge $e \in E_2$ with
 668 exactly one endpoint in $V(H')$.

17:16 A New Conjecture on 2-CSP's with Implications to Densest k -Subgraph

669 The running time of the algorithm is $O(T(n) \cdot \text{poly}(n))$.

670 Next, we obtain the following corollary.

671 **► Corollary 13.** *Assume that there exists a constant $0 < \varepsilon \leq 1/2$, and an $\alpha(N)$ -approximation*
 672 *algorithm \mathcal{A} for the Bipartite Densest (k_1, k_2) -Subgraph problem, whose running time is at*
 673 *most $T(N)$, where N is the number of vertices in the input graph, and $\alpha(N) = 2^{(4 \log N)^\varepsilon}$.*
 674 *Then there is an algorithm, whose input consists of an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$*
 675 *of Bipartite 2-CSP and parameter n that is greater than a large enough constant, so that*
 676 *size(\mathcal{I}) $\leq n$ holds, and \mathcal{I} is a $d(n)$ -to- $d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^\varepsilon}$.*
 677 *Let $\beta = 2^{8(\log n)^{1/2+\varepsilon}}$. The algorithm returns a partition (E_1, E_2, E_3) of $E(H(\mathcal{I}))$, where E_1*
 678 *is a bad set of constraints, E_2 is a β^3 -good set of constraints, and $|E_1 \cup E_2| \geq \frac{|E(H(\mathcal{I}))|}{\beta}$. The*
 679 *running time of the algorithm is $O(T(n) \cdot \text{poly}(n))$.*

680 **Proof:** Throughout the proof, we assume that there exists a constant $0 < \varepsilon \leq 1/2$, and
 681 an $\alpha(N)$ -approximation algorithm \mathcal{A} for the Bipartite Densest (k_1, k_2) -Subgraph problem,
 682 whose running time is at most $T(N)$, where N is the number of vertices in the input graph,
 683 and $\alpha(N) = 2^{(2 \log N)^\varepsilon}$. Assume that we are given an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of
 684 Bipartite 2-CSP, together with a parameter n that is greater than a large enough constant, so
 685 that size(\mathcal{I}) $\leq n$, and \mathcal{I} is a $d(n)$ -to- $d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^\varepsilon}$. For
 686 convenience, we denote $H = H(\mathcal{I})$. Our algorithm uses a parameter $\eta = 2^{12} d(n) \cdot \alpha(n) \cdot \log^4 n$.

687 The algorithm is iterative. Over the course of the algorithm, we maintain a collection \mathcal{H}
 688 of subgraphs of H , and another subgraph H^g of H . We will ensure that, throughout the
 689 algorithm, all graphs in $\mathcal{H} \cup \{H^g\}$ are mutually disjoint in their vertices. We denote by
 690 $E^g = E(H^g)$ and $E^1 = \bigcup_{H' \in \mathcal{H}} E(H')$. Additionally, we maintain another set E^b of edges of
 691 H , that is disjoint from $E^g \cup E^1$, and we denote by $E^0 = E(H) \setminus (E^g \cup E^b \cup E^1)$ the set
 692 of all remaining edges of H . We ensure that the following invariants hold throughout the
 693 algorithm.

- 694 **11.** set $E^g = E(H^g)$ of edges is β^3 -good;
 695 **12.** set E^b of edges is bad; and
 696 **13.** all graphs in $\mathcal{H} \cup \{H^g\}$ are disjoint in their vertices.

697 Intuitively, we will start with the set \mathcal{H} containing a single graph H , and $E^g = E^b =$
 698 $E^0 = \emptyset$. As the algorithm progresses, we will iteratively add edges to sets E^g, E^b and E^0 ,
 699 while partitioning the graphs in \mathcal{H} into smaller subgraphs. The algorithm will terminate
 700 once $\mathcal{H} = \emptyset$. The key in the analysis of the algorithm is to ensure that $|E^0|$ is relatively
 701 small when the algorithm terminates. We do so via a charging scheme: we assign a budget
 702 to every edge of $E^1 \cup E^g \cup E^b$, that evolves over the course of the algorithm, and we keep
 703 track of this budget over the course of the algorithm.

704 In order to define vertex budgets, we will assign, to every graph $H \in \mathcal{H}$ a *level*, that is an
 705 integer between 0 and $\lceil \log n \rceil$. We will ensure that, throughout the algorithm, the following
 706 additional invariants hold:

- 707 **14.** If $H' \in \mathcal{H}$ is a level- i graph, then the budget of every edge $e \in E(H')$ is at most η^i ; and
 708 **15.** Throughout the algorithm's execution, the total budget of all edges in $E^g \cup E^b \cup E^1$ is at
 709 least $|E(H)|$.

710 Intuitively, at the end of the algorithm, we will argue that the level of every graph in \mathcal{H}
 711 is not too large, and that the budget of every edge in $E^g \cup E^b \cup E^1$ is not too large. Since

712 the total budget of all edges in $E^g \cup E^b \cup E^1$ is at least $|E(H)|$, it will then follow that
 713 $|E^g \cup E^b \cup E^1|$ is sufficiently large. We now proceed to describe the algorithm.

714 Our algorithm will repeatedly use the algorithm from Corollary 12, with the same
 715 functions $\alpha(N)$, $d(n)$, and parameter β . In order to be able to use the corollary, we need
 716 to establish that $\beta \geq 2^{30}(\alpha(n))^3(\log n)^{12}$. This is immediate to verify since $\beta = 2^{8(\log n)^{1/2+\epsilon}}$,
 717 $\alpha(n) = 2^{(4 \log n)^\epsilon}$, and n is large enough.

718 At the beginning of the algorithm, we set $E^0 = E^g = E^b = \emptyset$, and we let \mathcal{H} contain a
 719 single graph H , which is assigned level 0. Note that $E^1 = E(H)$ must hold. Every edge
 720 $e \in E(H)$ is assigned budget $b(e) = 1$. Clearly, the total budget of all edges of $E^1 \cup E^g \cup E^b$
 721 is $B = \sum_{e \in E^1 \cup E^g \cup E^b} b(e) = |E(H)|$. The algorithm performs iterations, as long as $\mathcal{H} \neq \emptyset$.
 722 In every iteration, we select an arbitrary graph $H' \in \mathcal{H}$ to process.

723 We now describe an iteration where some graph $H' \in \mathcal{H}$ is processed. We assume
 724 that graph H' is assigned level i . Notice that graph H' naturally defines an instance
 725 $\mathcal{I}' = (X', Y', A, \mathcal{C}', \{G_C\}_{C \in \mathcal{C}'})$ of Bipartite 2-CSP, where $X' = V(H') \cap X$, $Y' = V(H') \cap Y$,
 726 $\mathcal{C}' = \{C \in \mathcal{C} \mid e_C \in E(H')\}$, and the graphs G_C for constraints $C \in \mathcal{C}'$ remain the same as
 727 in instance \mathcal{I} . Clearly, $\text{size}(\mathcal{I}') \leq \text{size}(\mathcal{I}) \leq n$, and $H(\mathcal{I}') = H'$. Furthermore, instance \mathcal{I}'
 728 remains a $d(n)$ -to- $d(n)$ instance. We apply the algorithm from Corollary 12 to instance \mathcal{I}' ,
 729 with parameters n and β remaining unchanged. Consider the partition (E_1, E_2) of $E(H')$
 730 that the algorithm returns. Recall that the set E_1 of edges is bad. We add the edges of E_1
 731 to set E^b . From Invariant I2 and Observation 7, set E^b of edges continues to be bad. If
 732 the algorithm from Corollary 12 certified that E_2 is a β^3 -good set of edges, then we update
 733 graph H^g to be $H^g \cup (H' \setminus E_1)$, and we add the edges of E_2 to set E^g . We then remove
 734 graph H' from \mathcal{H} , and continue to the next iteration. Note that, from Observation 9 and
 735 Invariants I1 and I3, the set E^g of edges continues to be β^3 -good. It is easy to verify that all
 736 remaining invariants also continue to hold.

737 From now on we assume that the algorithm from Corollary 12 returned a subgraph
 738 $H'' = (X'', Y'', E'')$ of H' , with $E'' \subseteq E_2$, such that $|X''| \leq \frac{2d(n) \cdot |X'|}{\beta}$ and $|Y''| \leq \frac{2d(n) \cdot |Y'|}{\beta}$.
 739 In particular, $|V(H'')| = |X''| + |Y''| \leq \frac{2d(n)}{\beta} \cdot (|X'| + |Y'|) \leq \frac{2d(n)}{\beta} \cdot |V(H')|$. Additionally,
 740 if we denote by E_2^* the subset of edges of E_2 containing all edges with exactly one endpoint
 741 in $X'' \cup Y''$, then $|E''| \geq \frac{|E_2^*|}{2048d(n) \cdot \alpha(n) \cdot \log^4 n}$ must hold. We let H^* be the graph obtained
 742 from $H' \setminus E_1$, by deleting the vertices of H'' from it, so $V(H^*) \cup V(H'') = V(H')$, and
 743 $E(H^*) \cup E(H'') \cup E_2^* = E_2$. We remove graph H' from \mathcal{H} , and we add graphs H'' and H^*
 744 to \mathcal{H} , with graph H'' assigned level $(i+1)$, and graph H^* assigned level i . We also add the
 745 edges of E_2^* to E^0 , and we update the set E^1 of edges to contain all edges of $\bigcup_{\tilde{H} \in \mathcal{H}} E(\tilde{H})$.
 746 Since we did not modify graph H^g in the current iteration, it is immediate to verify that
 747 Invariants I1–I3 continue to hold. Next, we update the budgets of edges, in order to ensure
 748 that Invariants I4 and I5 continue to hold. Intuitively, the edges of E_2^* are now added to
 749 set E^0 , so we need to distribute their budget among the edges of $E(H'')$, in order to ensure
 750 that the total budget of all edges in $E^g \cup E^b \cup E^1$ does not decrease. This will ensure that
 751 Invariant I5 continues to hold. At the same time, since the level of graph H'' is $(i+1)$, while
 752 the level of graph H' was i , we can increase the budgets of the edges of $E(H')$ and still
 753 maintain Invariant I4.

754 Formally, recall that Corollary 12 guarantees that $|E_2^*| \leq |E''| \cdot (2048d(n) \cdot \alpha(n) \cdot \log^4 n) =$
 755 $\frac{|E''| \cdot \eta}{2}$. From Invariant I4, the current budget of every edge in $E'' \cup E_2^*$ is bounded by η^i .
 756 Therefore, at the beginning of the current iteration: $\sum_{e \in E'' \cup E_2^*} b(e) \leq \eta^i \cdot (|E_2^*| + |E''|) \leq$
 757 $\eta^i \cdot |E''| \cdot (1 + \frac{\eta}{2}) < \eta^{i+1} \cdot |E''|$.

758 We set the budget of every edge in E'' to be η^{i+1} , and leave the budgets of all other
 759 edges unchanged. It is easy to verify that $\bigcup_{e \in E^g \cup E^b \cup E^1} b(e)$ does not decrease in the current

17:18 A New Conjecture on 2-CSP's with Implications to Densest k -Subgraph

760 iteration, so Invariant I5 continues to hold. It is also easy to verify that Invariant I4 continues
 761 to hold. Therefore, all invariants continue to hold at the end of the iteration. This completes
 762 the description of an iteration.

763 The algorithm terminates when $\mathcal{H} = \emptyset$. Clearly, we obtain a partition (E^g, E^b, E^0) of
 764 $E(H)$ into disjoint subsets, where the set E^b of edges is bad, and the set E^g of edges is
 765 β^3 -good. It remains to show that $|E^g \cup E^b| \geq \frac{|E(H)|}{\beta}$. We use the edge budgets in order to
 766 prove this. Let L^* be the largest level of any subgraph of H that belonged to \mathcal{H} at any time
 767 during the algorithm. We start with the following key observation, whose proof is deferred
 768 to the full version of the paper.

769 **► Observation 14.** $L^* \leq (\log n)^{1/2-\varepsilon}$.

770 From Invariant I4, throughout the algorithm, for every edge $e \in E^1$, $b(e) \leq \eta^{L^*}$ must
 771 hold. Once an edge is added to $E^b \cup E^g$, its budget does not change. Therefore, at the end
 772 of the algorithm, the budget of every edge in $E^g \cup E^b$ is at most η^{L^*} . On the other hand,
 773 from Invariant I5, at the end of the algorithm, the total budget of all edges in $E^1 \cup E^g \cup E^b$
 774 is at least $|E(H)|$. Therefore, at the end of the algorithm, $|E^g \cup E^b| \geq \frac{|E(H)|}{\eta^{L^*}}$ holds.

775 We now bound η^{L^*} . Recall that $\eta = 2^{12}d(n) \cdot \alpha(n) \cdot \log^4 n \leq 2^{4(\log n)^\varepsilon}$, since $d(n) \leq 2^{(\log n)^\varepsilon}$,
 776 $\alpha(n) = 2^{(4 \log n)^\varepsilon}$, and n is large enough. Since, from Observation 14, $L^* \leq (\log n)^{1/2-\varepsilon}$, we
 777 get that $\eta^{L^*} \leq 2^{4(\log n)^{1/2}} < \beta$, since $\beta = 2^{8(\log n)^{1/2+\varepsilon}}$. Therefore, $|E^g \cup E^b| \geq |E(H)|/\beta$ as
 778 required.

779 Lastly, it is easy to verify that the algorithm has at most $\text{poly}(n)$ iterations, and the
 780 running time of each iteration is bounded by $O(T(n) \cdot \text{poly}(n))$, so the total running time of
 781 the algorithm is at most $O(T(n) \cdot \text{poly}(n))$. \square

782 We are now ready to complete the proof of Theorem 10. Assume that there exists a
 783 constant $0 < \varepsilon \leq 1/2$, and an $\alpha(N)$ -approximation algorithm \mathcal{A} for the Densest k -Subgraph
 784 problem, whose running time is at most $T(N)$, where N is the number of vertices in the
 785 input graph, and $\alpha(N) = 2^{(\log N)^\varepsilon}$. From Lemma 2, there exists an $\alpha'(N)$ -approximation
 786 algorithm \mathcal{A} for the Bipartite Densest (k_1, k_2) -Subgraph problem, where N is the number of
 787 vertices in the input graph, and $\alpha'(N) \leq O(\alpha(N^2)) \leq O(2^{(2 \log N)^\varepsilon})$. The running time of
 788 the algorithm is at most $O(T(N^2) \cdot \text{poly}(N))$. Denote by $T'(N) = O(T(N^2) \cdot \text{poly}(N))$ this
 789 bound on the running time of the algorithm, and let $\alpha''(N) = 2^{(4 \log N)^\varepsilon}$. Then there is an
 790 $\alpha''(N)$ -approximation algorithm for Bipartite Densest (k_1, k_2) -Subgraph with running time at
 791 most $O(T'(N))$.

792 Assume now that we are given an instance $\mathcal{I} = (X, Y, A, \mathcal{C}, \{G_C\}_{C \in \mathcal{C}})$ of Bipartite 2-CSP
 793 and parameter n that is greater than a large enough constant, so that $\text{size}(\mathcal{I}) \leq n$ holds,
 794 and \mathcal{I} is a $d(n)$ -to- $d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^\varepsilon}$. Let $\beta = 2^{8(\log n)^{1/2+\varepsilon}}$,
 795 and let $r = \lceil \beta \cdot \log n \rceil$. For convenience, we denote $H = H(\mathcal{I})$. Initially, we set $E^b = \emptyset$. Our
 796 algorithm performs r iterations, where for all $1 \leq j \leq r$, in iteration j we construct the set
 797 $E_j \subseteq E(H)$ of edges, that is β^3 -good, and possibly adds some edges to set E^b . We ensure
 798 that, throughout the algorithm, the set E^b of edges is bad.

799 We now describe the j th iteration. We assume that sets E_1, \dots, E_{j-1} of edges of
 800 H were already defined. We construct graph H_j , that is obtained from graph H , by
 801 deleting the edges of $E_1 \cup \dots \cup E_{j-1} \cup E^b$ from it. Notice that graph H_j naturally defines
 802 an instance $\mathcal{I}_j = (X, Y, A, \mathcal{C}_j, \{G_C\}_{C \in \mathcal{C}_j})$ of Bipartite 2-CSP, with $H_j = H(\mathcal{I}_j)$, where
 803 $\mathcal{C}_j = \{C \in \mathcal{C} \mid e_C \in E(H_j)\}$. We apply the algorithm from Corollary 13 to graph H_j , with
 804 parameters n, β , and $d(n)$ remaining unchanged. Consider a partition (E^1, E^2, E^3) of $E(H_j)$
 805 that the algorithm returns. We add the edges of E^1 to set E^b . Since both sets of edges are
 806 bad, from Observation 7, set E^b of edges continues to be bad. We also set $E_j = E^2$, which

807 is guaranteed to be a β^3 -good set of edges. Recall that Corollary 13 also guarantees that
 808 $|E^1 \cup E^2| \geq |E(H_j)|/\beta$. We then continue to the next iteration.

809 Since, from the above discussion, for all $1 \leq j < r$, $|E(H_{j+1})| \leq \left(1 - \frac{1}{\beta}\right) |E(H_j)|$, and
 810 since $r = \lceil \beta \cdot \log n \rceil$, at the end of the algorithm, we are guaranteed that the final collection
 811 E^b, E_1, \dots, E_r of subsets of edges indeed partitions $E(H)$.

812 Notice that the running time of a single iteration is bounded by $O(T'(n) \cdot \text{poly}(n)) \leq$
 813 $O(T(\text{poly}(n)) \cdot \text{poly}(n))$. Since the number of iterations is bounded by $\text{poly}(n)$, the total
 814 running time of the algorithm is bounded by $O(T(\text{poly}(n)) \cdot \text{poly}(n))$.

815 **4** Reductions from Dense k -Coloring and (r,h) -Graph Partitioning to 816 Densest k -Subgraph

817 Our reductions from the Dense k -Coloring and (r,h) -Graph Partitioning problems to Densest
 818 k -Subgraph are summarized in the following theorem.

819 **► Theorem 15.** *Let $\alpha : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be an increasing function, such that $\alpha(n) \leq o(n)$. Assume*
 820 *that there is an efficient $\alpha(n)$ -approximation algorithm for the Densest k -Subgraph problem,*
 821 *where n is the number of vertices in the input graph. Then both of the following hold:*

- 822 **■** *there is an efficient randomized algorithm that, given an instance of Dense k -Coloring*
 823 *whose graph contains N vertices, with high probability computes an $O(\alpha(N^2) \cdot \text{poly} \log N)$ -*
 824 *approximate solution to this instance; and*
- 825 **■** *there is an efficient randomized algorithm that, given an instance of (r,h) -Graph Partitioning*
 826 *whose graph contains N vertices, with high probability computes an $O(\alpha(N^2) \cdot \text{poly} \log N)$ -*
 827 *approximate solution to this instance.*

828 The proof of the theorem is deferred to the full version of the paper, due to lack of space.
 829 We provide a high-level overview of the proof of the first assertion: a reduction from Dense
 830 k -Coloring to Densest k -Subgraph. The proof of the second assertion is similar. We start by
 831 considering an LP-relaxation of the Dense k -Coloring problem, whose number of variables
 832 is at least $\binom{N}{k}$. Due to this high number of variables, we cannot solve it directly. We first
 833 show an algorithm, that, given an approximate fractional solution to this LP-relaxation,
 834 whose support size is polynomial in N , computes an approximate integral solution to the
 835 Dense k -Coloring problem instance. We then show an efficient algorithm that computes an
 836 approximate solution to the LP-relaxation, whose support is relatively small. In order to do
 837 so, we design an approximate separation oracle to the dual LP of the LP-relaxation, that
 838 relies on an approximation algorithm for Densest k -Subgraph.

839 **5** Reductions from Densest k -Subgraph to Dense k -Coloring and 840 (r,h) -Graph Partitioning

841 Our reductions from Densest k -Subgraph to Dense k -Coloring and (r,h) -Graph Partitioning are
 842 summarized in the following theorem.

843 **► Theorem 16.** *Let $\alpha : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be an increasing function with $\alpha(n) \leq o(n)$. Then the*
 844 *following hold:*

- 845 **■** *If there exists an efficient $\alpha(n)$ -approximation algorithm A for the Dense k -Coloring prob-*
 846 *lem, where n is the number of vertices in the input graph, then there exists a randomized*
 847 *algorithm for the Densest k -Subgraph problem, whose running time is $N^{O(\log N)}$, that with*

848 *high probability computes an $O(\alpha(N^{O(\log N)}) \cdot \log N)$ -approximate solution to the input*
 849 *instance of the problem; here N is the number of vertices in the input instance of Densest*
 850 *k -Subgraph.*

851 ■ *If there exists an efficient $\alpha(n)$ -approximation algorithm for the (r,h)-Graph Partitioning*
 852 *problem, where n is the number of vertices in the input graph, then there exists a random-*
 853 *ized algorithm for the Densest k -Subgraph problem, whose running time is $N^{O(\log N)}$, that*
 854 *with high probability computes an $O((\alpha(N^{O(\log N)}))^3 \cdot \log^2 N)$ -approximate solution to the*
 855 *input instance of the problem; here N is the number of vertices in the input instance of*
 856 *Densest k -Subgraph.*

857 Due to lack of space, we defer the proof of the theorem to the full version of the paper.
 858 The key to both reductions is a randomized algorithm, that, given an instance $\text{DkS}(G, k)$ of
 859 the Densest k -Subgraph problem, constructs an auxiliary graph H . Intuitively, if instance
 860 $\text{DkS}(G, k)$ of Densest k -Subgraph has a solution of value h , then with high probability, graph
 861 H has close to $|V(H)|/k$ subgraphs that contain close to h edges each. On the other hand,
 862 there is an algorithm that, given a subgraph $H' \subseteq H$ that contains at most k vertices, extracts
 863 a subgraph of the original graph G , containing at most k vertices, and close to $|E(H')|$
 864 edges. If $|V(G)| = N$, then our construction of graph H ensures that $|V(H)| \leq N^{O(\log N)}$,
 865 which leads to the quasi-polynomial time of our reductions. The specific construction of the
 866 graph H is inspired by the ideas from [34]. We obtain the following immediate corollary of
 867 Theorem 16, whose proof is deferred to the full version of the paper due to lack of space.

868 ► **Corollary 17.** *Assume that Conjecture 3 holds and that $\text{NP} \not\subseteq \text{BPTIME}(n^{O(\log n)})$. Then*
 869 *for some constant $0 < \varepsilon' \leq 1/2$, there is no efficient $2^{(\log n)^{\varepsilon'}}$ -approximation algorithm for*
 870 *(r,h)-Graph Partitioning, and there is no efficient $2^{(\log n)^{\varepsilon'}}$ -approximation algorithm for Dense*
 871 *k -Coloring.*

872 **6** Reductions between (r,h)-Graph Partitioning and Maximum 873 Bounded-Crossing Subgraph

874 We establish a connection between the (r,h)-Graph Partitioning and Maximum Bounded-Crossing
 875 Subgraph problems via the following two theorems.

876 ► **Theorem 18.** *Let $\alpha : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be an increasing function with $\alpha(n) = o(n)$. Assume*
 877 *that there exists an efficient $\alpha(n)$ -approximation algorithm for the (r,h)-Graph Partitioning*
 878 *problem, where n is the number of vertices in the input graph. Then there exists an efficient*
 879 *$O(\alpha(N) \cdot \text{poly log } N)$ -approximation algorithm for Maximum Bounded-Crossing Subgraph,*
 880 *where N is the number of vertices in the input instance of Maximum Bounded-Crossing*
 881 *Subgraph.*

882 ► **Theorem 19.** *Let $\alpha : \mathbb{Z}^+ \rightarrow \mathbb{Z}^+$ be an increasing function with $\alpha(n) = o(n)$. Assume that*
 883 *there exists an efficient $\alpha(N)$ -approximation algorithm for the Maximum Bounded-Crossing*
 884 *Subgraph problem, where N is the number of vertices in the input graph. Then there exists an*
 885 *efficient $O((\alpha(n))^2 \cdot \text{poly log } n)$ -approximation algorithm for (r,h)-Graph Partitioning, where n*
 886 *is the number of vertices in the input instance of (r,h)-Graph Partitioning.*

887 The proofs of the above two theorems are deferred to the full version of the paper. Both
 888 proofs exploit well-known connections between crossing number and graph partitioning, that
 889 can be viewed as an extension of the classical Planar Separator Theorem of [45]: namely,
 890 if a graph G has a drawing with at most L crossings, then there is a balanced cut in G ,

891 containing at most $O\left(\sqrt{L + \Delta \cdot |E(G)|}\right)$ edges, where Δ is maximum vertex degree in G .
 892 Another useful fact exploited in the proofs of both these theorems is that any graph G with
 893 m edges has a plane drawing with at most m^2 crossings. For example, if $\mathcal{H} = \{H_1, \dots, H_r\}$
 894 is a solution to an instance of the (r, h) -Graph Partitioning problem on graph G , then there is a
 895 drawing of graph $H = \bigcup_{i=1}^r H_i$, in which the number of crossings is bounded by $r \cdot h^2$. These
 896 two facts are exploited in order to establish a close relationship between the (r, h) -Graph
 897 Partitioning and Maximum Bounded-Crossing Subgraph problems, and complete the proofs of
 898 Theorems 18 and 19. By combining Theorem 19 with Corollary 17, we obtain the following
 899 corollary, whose proof is deferred to the full version of the paper.

900 ► **Corollary 20.** *Assume that Conjecture 3 holds and that $\text{NP} \not\subseteq \text{BPTIME}(n^{O(\log n)})$. Then*
 901 *for some constant $0 < \varepsilon' \leq 1/2$, there is no efficient $2^{(\log n)^{\varepsilon'}}$ -approximation algorithm for*
 902 *Maximum Bounded-Crossing Subgraph.*

903 7 Acknowledgement

904 The authors thank Irit Dinur and Uri Feige for insightful and helpful discussions.

905 — References —

- 906 1 M. Ajtai, V. Chvátal, M. Newborn, and E. Szemerédi. Crossing-free subgraphs. *Theory and*
 907 *Practice of Combinatorics*, pages 9–12, 1982.
- 908 2 Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri Weinstein.
 909 Inapproximability of densest k -subgraph from average case hardness. *Manuscript*, 2011.
 910 <https://www.tau.ac.il/~nogaa/PDFS/dks8.pdf>.
- 911 3 Christoph Ambuhl, Monaldo Mastrolilli, and Ola Svensson. Inapproximability results for
 912 sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In *48th*
 913 *Annual IEEE Symposium on Foundations of Computer Science (FOCS'07)*, pages 329–337.
 914 IEEE, 2007.
- 915 4 Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games
 916 and related problems. *Journal of the ACM (JACM)*, 62(5):1–25, 2015.
- 917 5 Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David
 918 Steurer. Making the long code shorter. *SIAM Journal on Computing*, 44(5):1287–1324, 2015.
- 919 6 Boaz Barak, Pravesh K. Kothari, and David Steurer. Small-set expansion in shortcode graph
 920 and the 2-to-2 conjecture. In *10th Innovations in Theoretical Computer Science Conference,*
 921 *ITCS 2019, January 10–12, 2019, San Diego, California, USA*, pages 9:1–9:12, 2019. URL:
 922 <https://doi.org/10.4230/LIPIcs.ITCS.2019.9>, doi:10.4230/LIPIcs.ITCS.2019.9.
- 923 7 Siddharth Barman. Approximating nash equilibria and dense bipartite subgraphs via an
 924 approximate version of caratheodory’s theorem. In *Proceedings of the forty-seventh annual*
 925 *ACM symposium on Theory of computing*, pages 361–369, 2015.
- 926 8 Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan.
 927 Detecting high log-densities: an $O(n^{1/4})$ approximation for densest k -subgraph. In *Proceedings*
 928 *of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts,*
 929 *USA, 5–8 June 2010*, pages 201–210, 2010. URL: [http://doi.acm.org/10.1145/1806689.](http://doi.acm.org/10.1145/1806689.1806718)
 930 [1806718](http://doi.acm.org/10.1145/1806689.1806718), doi:10.1145/1806689.1806718.
- 931 9 Aditya Bhaskara, Moses Charikar, Venkatesan Guruswami, Aravindan Vijayaraghavan, and
 932 Yuan Zhou. Polynomial integrality gaps for strong sdp relaxations of densest k -subgraph. In
 933 *Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms*, pages
 934 388–405. SIAM, 2012.
- 935 10 Mark Braverman, Young Kun Ko, Aviad Rubinfeld, and Omri Weinstein. Eth hardness for
 936 densest- k -subgraph with perfect completeness. In *Proceedings of the Twenty-Eighth Annual*
 937 *ACM-SIAM Symposium on Discrete Algorithms*, pages 1326–1341. SIAM, 2017.

- 938 11 Sergio Cabello. Hardness of approximation for crossing number. *Discrete & Computational*
939 *Geometry*, 49(2):348–358, 2013.
- 940 12 Shih-Chia Chang, Li-Hsuan Chen, Ling-Ju Hung, Shih-Shun Kao, and Ralf Klasing. The
941 hardness and approximation of the densest k -subgraph problem in parameterized metric graphs.
942 In *2020 International Computer Symposium (ICS)*, pages 126–130. IEEE, 2020.
- 943 13 Chandra Chekuri and Anastasios Sidiropoulos. Approximation algorithms for euler genus
944 and related problems. In *2013 IEEE 54th Annual Symposium on Foundations of Computer*
945 *Science*, pages 167–176. IEEE, 2013.
- 946 14 Markus Chimani and Petr Hliněný. A tighter insertion-based approximation of the crossing
947 number. In *International Colloquium on Automata, Languages, and Programming*, pages
948 122–134. Springer, 2011.
- 949 15 Eden Chlamtác, Michael Dinitz, Christian Konrad, Guy Kortsarz, and George Rabanca. The
950 densest k -subhypergraph problem. *SIAM Journal on Discrete Mathematics*, 32(2):1458–1477,
951 2018.
- 952 16 Julia Chuzhoy. An algorithm for the graph crossing number problem. In *Proceedings of the*
953 *forty-third annual ACM symposium on Theory of computing*, pages 303–312. ACM, 2011.
- 954 17 Julia Chuzhoy. Excluded grid theorem: Improved and simplified. In *Proceedings of the*
955 *forty-seventh annual ACM symposium on Theory of Computing*, pages 645–654, 2015.
- 956 18 Julia Chuzhoy, David Hong Kyun Kim, and Rachit Nimavat. Almost polynomial hardness of
957 node-disjoint paths in grids. *Theory of Computing*, 17(6):1–57, 2021.
- 958 19 Julia Chuzhoy, Sepideh Mahabadi, and Zihan Tan. Towards better approximation of graph
959 crossing number. In *2020 IEEE 61st Annual Symposium on Foundations of Computer Science*
960 *(FOCS)*, pages 73–84. IEEE, 2020. Full version: Arxiv:2011.06545.
- 961 20 Julia Chuzhoy, Yury Makarychev, and Anastasios Sidiropoulos. On graph crossing number
962 and edge planarization. In *Proceedings of the twenty-second annual ACM-SIAM symposium*
963 *on Discrete algorithms*, pages 1050–1069. SIAM, 2011.
- 964 21 Julia Chuzhoy and Zihan Tan. A subpolynomial approximation algorithm for graph crossing
965 number in low-degree graphs. In *Proceedings of the 54th Annual ACM SIGACT Symposium*
966 *on Theory of Computing*, STOC 2022, pages 303–316, 2022. Full version: Arxiv:2202.06827.
- 967 22 Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expand-
968 ing sets in Grassmann graphs. In *Proceedings of the 50th Annual ACM SIGACT Symposium on*
969 *Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 940–951,
970 2018. URL: <https://doi.org/10.1145/3188745.3188806>, doi:10.1145/3188745.3188806.
- 971 23 Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the
972 2-to-1 games conjecture? In *Proceedings of the 50th Annual ACM SIGACT Symposium on*
973 *Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018*, pages 376–389,
974 2018. URL: <https://doi.org/10.1145/3188745.3188804>, doi:10.1145/3188745.3188804.
- 975 24 Guy Even, Sudipto Guha, and Baruch Schieber. Improved approximations of crossings in
976 graph drawings and vlsi layout areas. *SIAM Journal on Computing*, 32(1):231–252, 2002.
- 977 25 Uriel Feige. Relations between average case complexity and approximation complexity. In
978 *Proceedings of the thirty-fourth annual ACM symposium on Theory of computing*, pages 534–543,
979 2002.
- 980 26 Uriel Feige and Michael Langberg. Approximation algorithms for maximization problems
981 arising in graph partitioning. *Journal of Algorithms*, 41(2):174–211, 2001.
- 982 27 Uriel Feige, David Peleg, and Guy Kortsarz. The dense k -subgraph problem. *Algorithmica*,
983 29(3):410–421, 2001.
- 984 28 Uriel Feige, Michael Seltser, et al. On the densest k -subgraph problem. Technical Report
985 CS97-16, Weizmann Institute of Science., 1997. [https://citeseerx.ist.psu.edu/viewdoc/
986 download?doi=10.1.1.37.9962&rep=rep1&type=pdf](https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.37.9962&rep=rep1&type=pdf).
- 987 29 Doron Goldstein and Michael Langberg. The dense k subgraph problem. *arXiv preprint*
988 *arXiv:0912.5327*, 2009.

- 989 30 Tesshu Hanaka. Computing densest k -subgraph with structural parameters. *arXiv preprint*
990 *arXiv:2207.09803*, 2022.
- 991 31 Johan Håstad. Some optimal inapproximability results. *Journal of the ACM (JACM)*,
992 48(4):798–859, 2001.
- 993 32 Ken-ichi Kawarabayashi and Anastasios Sidiropoulos. Polylogarithmic approximation for
994 minimum planarization (almost). In *58th IEEE Annual Symposium on Foundations of*
995 *Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017*, pages 779–788,
996 2017. URL: <https://doi.org/10.1109/FOCS.2017.77>, doi:10.1109/FOCS.2017.77.
- 997 33 Ken-ichi Kawarabayashi and Anastasios Sidiropoulos. Polylogarithmic approximation for euler
998 genus on bounded degree graphs. In *Proceedings of the 51st Annual ACM SIGACT Symposium*
999 *on Theory of Computing*, pages 164–175. ACM, 2019.
- 1000 34 Sanjeev Khanna, Nathan Linial, and Shmuel Safra. On the hardness of approximating the
1001 chromatic number. *Combinatorica*, 20(3):393–415, 2000.
- 1002 35 Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the*
1003 *thirty-fourth annual ACM symposium on Theory of computing*, pages 767–775, 2002.
- 1004 36 Subhash Khot. Ruling out ptas for graph min-bisection, dense k -subgraph, and bipartite
1005 clique. *SIAM Journal on Computing*, 36(4):1025–1071, 2006.
- 1006 37 Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra. Small set expansion in the
1007 Johnson graph. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:78, 2018.
1008 <https://eccc.weizmann.ac.il/report/2018/078>.
- 1009 38 Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and Grassmann
1010 graphs. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing,*
1011 *STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 576–589, 2017. URL: <https://doi.org/10.1145/3055399.3055432>, doi:10.1145/3055399.3055432.
- 1012 39 Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in Grassmann graph have
1013 near-perfect expansion. In *59th IEEE Annual Symposium on Foundations of Computer*
1014 *Science, FOCS 2018, Paris, France, October 7-9, 2018*, pages 592–601, 2018. URL: <https://doi.org/10.1109/FOCS.2018.00062>, doi:10.1109/FOCS.2018.00062.
- 1015 40 Subhash Khot and Muli Safra. A two-prover one-round game with strong soundness. *Theory*
1016 *of Computing*, 9:863–887, 2013. URL: <https://doi.org/10.4086/toc.2013.v009a028>, doi:
1017 10.4086/toc.2013.v009a028.
- 1018 41 G Kortsarz and D Peleg. On choosing a dense subgraph. In *Proceedings of 1993 IEEE 34th*
1019 *Annual Foundations of Computer Science*, pages 692–701. IEEE Computer Society, 1993.
- 1020 42 F. T. Leighton. *Complexity issues in VLSI: optimal layouts for the shuffle-exchange graph and*
1021 *other networks*. MIT Press, 1983.
- 1022 43 Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in
1023 designing approximation algorithms. *Journal of the ACM (JACM)*, 46(6):787–832, 1999.
- 1024 44 Bingkai Lin. The parameterized complexity of the k -biclique problem. *Journal of the ACM*
1025 *(JACM)*, 65(5):1–23, 2018.
- 1026 45 Richard J Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. *SIAM*
1027 *Journal on Applied Mathematics*, 36(2):177–189, 1979.
- 1028 46 Pasin Manurangsi. Almost-polynomial ratio ETH-hardness of approximating densest k -
1029 subgraph. In *Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of*
1030 *Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017*, pages 954–961, 2017. URL:
1031 <http://doi.acm.org/10.1145/3055399.3055412>, doi:10.1145/3055399.3055412.
- 1032 47 Pasin Manurangsi. Inapproximability of maximum biclique problems, minimum k -cut and
1033 densest at-least- k -subgraph from the small set expansion hypothesis. *Algorithms*, 11(1):10,
1034 2018.
- 1035 48 J. Matoušek. *Lectures on discrete geometry*. Springer-Verlag, 2002.
- 1036 49 J. Pach and G. Tóth. Thirteen problems on crossing numbers. *Geombinatorics*, 9(4):194–207,
1037 2000.

17:24 A New Conjecture on 2-CSP's with Implications to Densest k -Subgraph

- 1040 50 R. B. Richter and G. Salazar. Crossing numbers. In L. W. Beineke and R. J. Wilson, editors,
1041 *Topics in Topological Graph Theory*, chapter 7, pages 133–150. Cambridge University Press,
1042 2009.
- 1043 51 Marcus Schaefer. The graph crossing number and its variants: A survey. *The electronic journal*
1044 *of combinatorics*, pages DS21–Sep, 2012.
- 1045 52 Renata Sotirov. On solving the densest k -subgraph problem on large graphs. *Optimization*
1046 *Methods and Software*, 35(6):1160–1178, 2020.
- 1047 53 David Steurer. Subexponential algorithms for d -to-1 two-prover games and for certifying
1048 almost perfect expansion. Available at [https://citeseerx.ist.psu.edu/viewdoc/download?](https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.189.5388&rep=rep1&type=pdf)
1049 [doi=10.1.1.189.5388&rep=rep1&type=pdf](https://citeseerx.ist.psu.edu/viewdoc/download?doi=10.1.1.189.5388&rep=rep1&type=pdf), 2010.
- 1050 54 P. Turán. A note of welcome. *J. Graph Theory*, 1:1–5, 1977.