# A New Conjecture on Hardness of Low-Degree 2-CSPs with Implications to Hardness of Densest $k$-Subgraph and Other Problems 

Julia Chuzhoy* Mina Dalirrooyfard ${ }^{\dagger} \quad$ Vadim Grinberg ${ }^{\ddagger} \quad$ Zihan Tan ${ }^{\S}$

November 14, 2022


#### Abstract

We propose a new conjecture on hardness of low-degree 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 further 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.


[^0]
## Contents

1 Introduction ..... 1
1.1 A More Detailed Overview of our Results and Techniques ..... 6
2 Preliminaries ..... 9
2.1 General Notation ..... 10
2.2 Problem Definitions and Additional Notation ..... 10
2.3 Chernoff Bound ..... 11
2.4 Auxiliary Lemma ..... 12
3 Conditional Hardness of Densest $k$-Subgraph ..... 12
3.1 Low-Degree CSP Conjecture ..... 12
3.2 Conditional Hardness of Densest $k$-Subgraph ..... 13
3.3 Proof of Theorem 3.7 ..... 15
3.4 Proof of Theorem 3.8 ..... 21
3.4.1 Step 1: Regularization ..... 22
3.4.2 Step 2: Assignment Graph and Reduction to Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph ..... 23
3.4.3 Step 3: Further Regularization ..... 24
3.4.4 $\quad$ Step 4: Certifying that $H$ is a Good Graph or Computing a Subgraph of $H$ ..... 27
4 Reductions from Dense $k$-Coloring and (r,h)-Graph Partitioning to Densest $k$-Subgraph ..... 29
4.1 An LP-Relaxation and Its Rounding ..... 30
4.2 Approximately Solving the LP-Relaxation ..... 33
5 Reductions from Densest $k$-Subgraph to Dense $k$-Coloring and (r,h)-Graph Partitioning ..... 43
5.1 Auxiliary Graph $H$ ..... 44
5.2 Completing the Reduction from Densest $k$-Subgraph to Dense $k$-Coloring ..... 49
5.3 Completing the Reduction from Densest $k$-Subgraph to (r,h)-Graph Partitioning ..... 50
6 Reductions between (r,h)-Graph Partitioning and Maximum Bounded-Crossing Subgraph ..... 52
6.1 Auxiliary Lemmas ..... 53
6.2 Reduction from Maximum Bounded-Crossing Subgraph to (r,h)-Graph Partitioning:
Proof of Theorem|6.1 ..... 56
6.3 Reduction from (r,h)-Graph Partitioning to Maximum Bounded-Crossing Subgraph- Proof of Theorem|6.2. ..... 58
6.3.1 Case 1: $C^{*} \geq 16 n \alpha(n) \log ^{7} n$ ..... 59
6.3.2 Case 2: $C^{*}<16 n \alpha(n) \log ^{7} n$ ..... 61
7 Acknowledgement ..... 64
A Proof of Lemma 2.1 ..... 65
A. 1 Reduction from Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph to Densest $k$-Subgraph ..... 65
A. 2 Reduction from Densest $k$-Subgraph to Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph ..... 66
B Reduction from (r,h)-Graph Partitioning to Densest $k$-Subgraph ..... 67
B. 1 Linear Programming Relaxation and an LP-Rounding Algortihm ..... 67
B. 2 Approximately Solving the LP-Relaxation ..... 69

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

We prove a new conditional hardness of approximation result for Densest $k$-Subgraph based on Low-Degree CSP Conjecture. In addition to the Densest $k$-Subgraph problem, we study three other problems. The first problem, called ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, recently arose in the hardness of approximation proof of the Node-Disjoint Paths problem of CKN21, who mention that the problem appears similar to Densest $k$-Subgraph, but could not formalize this intuition. We also study a new problem that we call Dense $k$-Coloring, that can be viewed as a natural middle ground between Densest $k$-Subgraph and ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning. The fourth problem that we study is a variation of the notoriously difficult Minimum Crossing Number problem, that we call Maximum BoundedCrossing Subgraph. This problem also arose implicitly in CKN21. We show that all four problems are roughly equivalent from the approximation perspective, in the regime where the approximation factors are somewhat large (but some of our reductions require quasi-polynomial time). We then derive conditional hardness of approximation results for all these problems based on these reductions and the conditional hardness of Densest $k$-Subgraph.
The main contribution of this paper is thus twofold: first, we propose a new conjecture on hardness of CSP's and show that a number of interesting hardness of approximation results follow from it. Second, we establish a close connection between the four problems that we study. The remainder of the Introduction is organized as follows. We start by providing a brief overview of the four problems that we study in this paper. We then state the Low-Degree CSP Conjecture and put it into context with existing results and well-known conjectures. Finally, we provide a more detailed overview of our results and techniques.

Densest $k$-Subgraph. In the Densest $k$-Subgraph problem, given an $n$-vertex graph $G$ and an integer $k>1$, the goal is to compute a subset $S$ of $k$ vertices of $G$, while maximizing the number of edges in $G[S]$. Densest $k$-Subgraph is one of the most basic graph optimization problems that has been studied extensively (see e.g. KP93, FS ${ }^{+} 97$, FPK01, FL01, Fei02, Kho06, GL09, BCC ${ }^{+} 10$, $\mathrm{AAM}^{+} 11, \mathrm{BCG}^{+} 12, \mathrm{Bar} 15, \mathrm{BKRW17}, \mathrm{Man} 17, \mathrm{CDK}^{+} 18, \mathrm{Man} 18, \mathrm{Lin} 18$, Sot20, $\left.\mathrm{CCH}^{+} 20, \mathrm{Han} 22 \mid\right)$. At the same time it seems notoriously difficult, and despite this extensive work, our understanding of its approximability is still incomplete. The best current approximation algorithm for Densest $k$-Subgraph, due to $\mathrm{BCC}^{+} 10$, achieves, for every $\varepsilon>0$, an $O\left(n^{1 / 4+\varepsilon}\right)$-approximation, in time
$n^{O(1 / \varepsilon)}$. Even though the problem appears to be very hard, its hardness of approximation proof has been elusive. For example, no constant-factor hardness of approximation proofs for Densest $k$-Subgraph are currently known under the standard $\mathrm{P} \neq \mathrm{NP}$ assumption, or even the stronger assumption that NP $\nsubseteq \operatorname{BPTIME}\left(n^{\text {poly } \log n}\right)$. In a breakthrough result, Khot Kho06 proved a factor- $c$ hardness of approximation for Densest $k$-Subgraph, for some small constant $c$, assuming that NP $\nsubseteq \cap_{\varepsilon>0} \operatorname{BPTIME}\left(2^{n^{\varepsilon}}\right)$. Several other papers proved constant and super-constant hardness of approximation results for Densest $k$-Subgraph under average-case complexity assumptions: namely that no efficient algorithm can refute random 3-SAT or random $k$-AND formulas [Fei02, $\mathrm{AAM}^{+} 11$ ]. Additionally, a factor $2^{\Omega\left(\log ^{2 / 3} n\right)}$-hardness of approximation was shown under assumptions on solving Planted Clique $\mathrm{AAM}^{+} 11$. In a recent breakthrough, Manurangsi Man17 proved that, under the Exponential Time Hypothesis (ETH), the Densest $k$-Subgraph problem is hard to approximate to within factor $n^{1 /(\log \log n)^{c}}$, for some constant $c$. Proving a super-constant hardness of Densest $k$-Subgraph under weaker complexity assumptions remains a tantalizing open question that we attempt to address in this paper. Unfortunately, it seems unlikely that the techniques of Man17 can yield such a result. In this paper we show that, assuming the Low-Degree CSP Conjecture that we introduce, Densest $k$-Subgraph is NP-hard to approximate to within factor $2^{(\log n)^{\varepsilon}}$, for some constant $\varepsilon>0$.

The $(r, h)$-Graph Partitioning Problem. A recent paper CKN21 on the hardness of approximation of the Node-Disjoint Paths (NDP) problem formulated and studied a new graph partitioning problem, called ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning. The input to the problem is a graph $G$, and two integers, $r$ and $h$. The goal is to compute $r$ vertex-disjoint subgraphs $H_{1}, \ldots, H_{r}$ of $G$, such that for each $1 \leq i \leq r,\left|E\left(H_{i}\right)\right| \leq h$, while maximizing $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right|$. A convenient intuitive way of thinking about this problem is that we are interested in obtaining a balanced partition of the graph $G$ into $r$ vertex-disjoint subgraphs, so that the subgraphs contain sufficiently many edges. Unlike standard graph partitioning problems, that typically aim to minimize the number of edges connecting the different subgraphs in the solution, our goal is to maximize the total number of edges that are contained in the subgraphs. In order to avoid trivial solutions, in which one of the subgraphs contains almost the entire graph $G$, and the remaining subgraphs are almost empty, we place an upper bound $h$ on the number of edges that each subgraph may contribute towards the solution. Note that the subgraphs $H_{i}$ of $G$ in the solution need not be vertex-induced subgraphs.

The work of CKN21 attempted to use ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning as a proxy problem for proving hardness of approximation of NDP. Their results imply that NDP is at least as hard to approximate as ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, to within polylogarithmic factors. In order to prove hardness of NDP, it would then be sufficient to show that (r,h)-Graph Partitioning is hard to approximate. Unfortunately, CKN21 were unable to do so. Instead, they considered a generalization of ( $r, h$ )-Graph Partitioning, called ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning with Bundles. They showed that NDP is at least as hard as ( $\mathrm{r}, \mathrm{h}$ )Graph Partitioning with Bundles, and then proved hardness of this new problem. In the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning with Bundles problem, the input is the same as in ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, but now graph $G$ must be bipartite, and, for every vertex $v$, we are given a partition $\mathcal{B}(v)$ of the set of edges incident to $v$ into subsets that are called bundles. We require that, in a solution $\left(H_{1}, \ldots, H_{r}\right)$ to the problem, for every vertex $v \in V(G)$, and every bundle $\beta \in \mathcal{B}(v)$, at most one edge of $\beta$ contributes to the solution; in other words, at most one edge of $\beta$ may lie in $\bigcup_{i} E\left(H_{i}\right)$. This is a somewhat artificial problem, but this definition allows one to bypass some of the barriers that arise when
trying to prove the hardness of $(r, h)$-Graph Partitioning from existing hardness results for CSP's.
It was noted in CKN21 that the ( $r$,h)-Graph Partitioning problem resembles the Densest $k$-Subgraph problem for two reasons. First, in Densest $k$-Subgraph, the goal is to compute a dense subgraph of a given graph, with a prescribed number of vertices. One can think of ( $r, h$ )-Graph Partitioning as the problem of computing many vertex-disjoint dense subgraphs of a given graph. Second, natural hardness of approximation proofs for both problems seem to run into the same barriers. It is therefore natural to ask: (i) Can we prove that the ( $r, h$ )-Graph Partitioning problem itself is hard to approximate? In particular, can the techniques of [CKN21] be exploited in order to obtain such a proof? and (ii) Can we formalize this intuitive connection between ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning and Densest $k$-Subgraph? In this paper we make progress on both these questions. Our conditional hardness result for Densest $k$-Subgraph indeed builds on the ideas from CKN21 for proving hardness of ( $r, h$ )-Graph Partitioning with Bundles. We also provide "almost" approximation-preserving reductions between ( $r, h$ )-Graph Partitioning to Densest $k$-Subgraph: we show that, if there is an efficient factor $\alpha(n)$-approximation algorithm for Densest $k$-Subgraph, then there is a randomized efficient factor $O\left(\alpha\left(n^{2}\right) \cdot\right.$ poly $\left.\log n\right)$-approximation algorithm to (r,h)-Graph Partitioning. We also provide a reduction in the opposite direction: we prove that, if there is an efficient $\alpha(n)$-approximation algorithm for ( $r, h$ )-Graph Partitioning, then there is a randomized algorithm for Densest $k$-Subgraph, that achieves approximation factor $O\left(\left(\alpha\left(n^{O(\log n)}\right)\right)^{3} \cdot \log ^{2} n\right)$, in time $n^{O(\log n)}$. Therefore, we prove that Densest $k$-Subgraph and $(r, h)$-Graph Partitioning are roughly equivalent from the approximation perspective (at least for large approximation factors and quasi-polynomial running times). Combined with our conditional hardness of approximation for Densest $k$-Subgraph, our results show that, assuming the Low-Degree CSP Conjecture, for some constant $0<\varepsilon \leq 1 / 2$, there is no efficient $2^{(\log n)^{\varepsilon}}$-approximation algorithm for $(r, h)$-Graph Partitioning, unless NP $\subseteq \operatorname{BPTIME}\left(n^{O(\log n)}\right)$.

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

The Minimum Crossing Number problem was initially introduced by Turán Tur77 in 1944, and has been extensively studied since then (see, e.g., Chu11, CMS11, CH11, CS13, KS17, KS19, CMT20, and also RS09, PT00, Mat02, Sch12] for excellent surveys). But despite all this work, most aspects of the problem are still poorly understood. A long line of work LR99, EGS02, CMS11, Chu11, KS17, KS19, Chu15, CT22 has recently led to the first sub-polynomial approximation algorithm for the problem in low degree graphs. Specifically, CT22 obtain a factor $O\left(2^{O\left((\log n)^{7 / 8} \log \log n\right)} \cdot \Delta^{O(1)}\right)$ approximation algorithm for Minimum Crossing Number, where $\Delta$ is the maximum vertex degree. To the best of our knowledge, no non-trivial approximation algorithms are known for the problem when vertex degrees in the input graph $G$ can be arbitrary. However, on the negative side, only APX-hardness is known for the problem Cab13, AMS07]. As the current understanding of the Minimum Crossing Number problem from the approximation perspective is extremely poor, it is natural to study hardness of approximation of its variants.

Let us consider two extreme variations of the Minimum Crossing Number problem. The first variant
is the Minimum Crossing Number problem itself, where we need to draw an input graph $G$ in the plane with fewest crossings. The second variant is where we need to compute a subgraph $G^{\prime}$ of the input graph $G$ that is planar, while maximizing $\left|E\left(G^{\prime}\right)\right|$. The latter problem has a simple constant-factor approximation algorithm, obtained by letting $G^{\prime}$ be any spanning forest of $G$ (this is since a planar $n$-vertex graph may only have $O(n)$ edges).
In this paper we study a variation of the Minimum Crossing Number problem, that we call Maximum Bounded-Crossing Subgraph, which can be viewed as an intermediate problem between these two extremes. In the Maximum Bounded-Crossing Subgraph problem, given an $n$-vertex graph $G$ and an integer $L>0$, the goal is to compute a subgraph $H \subseteq G$, such that $H$ has a plane drawing with at most $L$ crossings, while maximizing $|E(H)|$. This problem is only interesting when the bound $L$ on the number of crossings is $\Omega(n)$. This is since, from the Crossing Number Inequality ACNS82 Lei83, if $|E(G)| \geq 4|V(G)|$, then the crossing number of $G$ is at least $\Omega\left(|E(G)|^{3} /|V(G)|^{2}\right)$. Therefore, for $L=O(n)$, a spanning tree provides a constant-factor approximation to the problem. We emphasize that the focus here is on dense graphs, whose crossing number may be as large as $\Omega\left(n^{4}\right)$.

The Maximum Bounded-Crossing Subgraph problem was implicitly used in CKN21 for proving hardness of approximation of NDP, as an intermediate problem, in the reduction from ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning with Bundles to NDP. Their work suggests that there may be a connection between ( $\mathrm{r}, \mathrm{h}$ )Graph Partitioning and Maximum Bounded-Crossing Subgraph, even though the two problems appear quite different. In this paper we prove that the two problems are roughly equivalent from the approximation perspective: if there is an efficient factor $\alpha(n)$-approximation algorithm for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, then there is an efficient $O(\alpha(n) \cdot$ poly $\log n)$-approximation algorithm for Maximum Bounded-Crossing Subgraph. On the other hand, an efficient $\alpha(n)$-approximation algorithm for Maximum Bounded-Crossing Subgraph implies an efficient $O\left((\alpha(n))^{2} \cdot\right.$ poly $\left.\log n\right)$-approximation algorithm for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning. Combined with our conditional hardness of approximation for (r,h)-Graph Partitioning, we get that, assuming the Low-Degree CSP Conjecture, for some constant $0<\varepsilon \leq 1 / 2$ there is no efficient $2^{(\log n)^{\varepsilon}}$-approximation algorithm for Maximum Bounded-Crossing Subgraph, unless NP $\subseteq \operatorname{BPTIME}\left(n^{O(\log n)}\right)$.

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

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

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

The famous Unique-Games Conjecture of Khot Kho02] applies to 1-to-1 CSP's. The conjecture states that, for any $0<\varepsilon<1$, there is a large enough value $A$, such that the ( $1-\varepsilon, \varepsilon$ )-Gap-CSP problem is NP-hard for 1-to-1 instances with alphabet size $A$. The conjecture currently remains open, though interesting progress has been made on the algorithmic side: the results of ABS15 provide an algorithm for the problem with running time $2^{n^{O\left(1 / \varepsilon^{1 / 3}\right)}}$.
A conjecture that is closely related to the Unique-Games Conjecture is the $d$-to- 1 Conjecture of Khot Kho02]. The conjecture states that, for every $0<\varepsilon<1$, and $d>0$, there is a large enough value $A$, such that the $(1, \varepsilon)$-Gap-CSP problem in $d$-to- 1 instances with alphabet size $A$ is NP-hard.

Håstad [Hås01] proved the following nearly optimal hardness of approximation results for CSP's: he showed that for every $0<\varepsilon<1$, there are values $d$ and $A$, such that the problem of $(1, \varepsilon)$-Gap-CSP in $d$-to- 1 instances with alphabet size $A$ is NP-hard. The value $d$, however, depends exponentially on poly $(1 / \varepsilon)$ in this result. In contrast, in the $d$-to- 1 Conjecture, both $d$ and $\varepsilon$ are fixed, and $d$ may not have such a strong dependence on $1 / \varepsilon$.
On the algorithmic side, the results of ABS15, Ste provide an algorithm for $(c, s)$-Gap-CSP on $d$-to- 1 instances. The running time of the algorithm is $2^{n^{O\left(1 /(\log (1 / s))^{1 / 2}\right)}}$, where the $O(\cdot)$ notation hides factors that are polynomial in $d$ and $A$.
A recent breakthrough in this area is the proof of the 2-to-2 conjecture (now theorem), that builds on a long sequence of work $\mathrm{BGH}^{+} 15, \mathrm{KS13}$ KMMS18, BKS19, KMS17, $\mathrm{DKK}^{+} 18 \mathrm{a}, \mathrm{DKK}^{+}$18b, KMS18.

The theorem proves that for every $0<\varepsilon<1$, there is a large enough value $A$, such that the ( $1-\varepsilon, \varepsilon$ )-Gap-CSP problem is NP-hard on 2-to-2 instances with alphabet size $A$.

In this paper, we propose the following conjecture, that we refer to as Low-Degree CSP Conjecture, regarding the hardness of Gap-CSP in $d$-to- $d$ instances.

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

We now compare this conjecture to existing conjectures and results in this area that we are aware of. First, in contrast to the $d$-to- 1 conjecture, we allow the parameter $d$ and the soundness parameter $s$ to be functions of $n$ - the size of the input instance. Note that the size of the input instance depends on the alphabet size $A$, so, unlike in the setting of the $d$-to- 1 conjecture, $A$ may no longer be arbitrarily large compared to $d$ and $s$.
The hardness of approximation result of Håstad Hås01 for $d$-to- $d$ CSP's only holds when $d$ depends exponentially on poly $(1 / s)$, (in particular it may not extend to the setting where $s(n)=$ $1 / 2^{64(\log n)^{1 / 2+\varepsilon}}$, since the size $n$ of the instance depends polynomially on $\left.d(n)\right)$.
We can also combine standard constant hardness of approximation results for CSP's (such as, for example, 3-SAT) with the Parallel Repetition theorem, to obtain NP-hardness of ( $1, s(n)$ )-Gap-CSP on $d(n)$-to- $d(n)$ instances. Using this approach, if we start from an instance of CSP of size $N$ and a constant hardness gap (with perfect completeness), after $\ell$ rounds of parallel repetition, we obtain hardness of $(1, s)$-Gap-CSP on $d$-to- $d$ instances with $s=2^{O(\ell)}, d=2^{O(\ell)}$, and the resulting instance size $n=N^{O(\ell)}$. By setting the number of repetition to be $\ell=\Theta\left((\log N)^{(1 / 2+\varepsilon) /(1 / 2-\varepsilon)}\right)$, we can ensure the desired bound $s(n)=1 / 2^{64(\log n)^{1 / 2+\varepsilon}}$. However, in this setting, we also get that $d(n)=2^{\Omega\left((\log n)^{1 / 2+\varepsilon}\right)}$, which is significantly higher than the desired value $d(n)=2^{(\log n)^{\varepsilon}}$.
Lastly, one could attempt to combine the recent proof of the 2-to-2 conjecture with Parallel Repetition in order to reap the benefits of both approaches, but the resulting parameters also fall short of the ones stated in the conjecture.
From the above discussion, one can view the Low-Degree CSP Conjecture as occupying a middle ground between the $d$-to-1 conjecture, and the results one can obtain via standard techniques of amplifying a constant hardness of a CSP, such as 3SAT, via Parallel Repetition.

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

### 1.1 A More Detailed Overview of our Results and Techniques

In addition to posing the Low-Degree CSP Conjecture that we already described above, we prove conditional hardness of approximation of the four problems that we consider. We also prove that all four problems are roughly equivalent approximation-wise. We now discuss the conditional hardness of approximation for Densest $k$-Subgraph and the connections between the four problems that we establish.

Conditional Hardness of Densest $k$-Subgraph. Our first result is a conditional hardness of Densest $k$-Subgraph. Specifically, we prove that, assuming Conjecture 1 holds and $\mathrm{P} \neq \mathrm{NP}$, for some $0<\varepsilon \leq 1 / 2$, there is no efficient approximation algorithm for Densest $k$-Subgraph problem that achieves approximation factor $2^{(\log N)^{\varepsilon}}$, where $N$ is the number of vertices in the input graph.
We now provide a brief overview of our techniques. The proof of the above result employs a Cook-type reduction, and follows some of the ideas that were introduced in CKN21. We assume for contradiction that there is a factor- $\alpha$ algorithm $\mathcal{A}$ for the Densest $k$-Subgraph problem, where $\alpha=2^{(\log N)^{\varepsilon}}$. Given an input instance $\mathcal{I}$ of the 2 -CSP problem of size $n$, that is a $d(n)$-to- $d(n)$ instance, we construct a constraint graph $H$ representing $\mathcal{I}$. We gradually decompose graph $H$ into a collection $\mathcal{H}$ of disjoint subgraphs, such that, for each subgraph $H^{\prime} \in \mathcal{H}$, we can either certify that the value of the corresponding instance of 2 - CSP is at most $1 / 4$, or it is at least $\beta$, for some carefully chosen parameter $\beta$. In order to compute the decomposition, we start with $\mathcal{H}=\{H\}$. If, for a graph $H^{\prime} \in \mathcal{H}$, we certified that the corresponding instance of 2-CSP has value at most $1 / 4$, or at least $\beta$, then we say that graph $H^{\prime}$ is inactive. Otherwise, we say that it is active. As long as $\mathcal{H}$ contains at least one active graph, we perform iterations. In each iteration, we select an arbitrary graph $H^{\prime} \in \mathcal{H}$ to process. In order to process $H^{\prime}$, we consider an assignment graph $G^{\prime}$ associated with $H^{\prime}$, that contains a vertex for every variable-assignment pair $(x, a)$, where $x$ is a variable whose corresponding vertex belongs to $H^{\prime}$. We view $G^{\prime}$ as an instance of the Densest $k$-Subgraph problem, for an appropriately chosen parameter $k$, and apply the approximation algorithm $\mathcal{A}$ for Densest $k$-Subgraph to it. Let $S$ be the set of vertices of $G^{\prime}$ that Algorithm $\mathcal{A}$ computes as a solution to this instance. We exploit the set $S$ of vertices in order to either (i) compute a large subset $E^{\prime} \subseteq E\left(H^{\prime}\right)$ of edges, such that, if we denote by $\mathcal{C}^{\prime} \subseteq \mathcal{C}$ the set of constraints corresponding to $E^{\prime}$, then at most $1 / 4$ of the constraints of $\mathcal{C}^{\prime}$ can be simultaneously satisfied; or (ii) compute a large subset $E^{\prime} \subseteq E\left(H^{\prime}\right)$ of edges as above, and certify that at least a $\beta$-fraction of such constraints can be satisfied; or (iii) compute a subgraph $H^{\prime \prime} \subseteq H^{\prime}$, such that $\left|V\left(H^{\prime \prime}\right)\right| \ll\left|V\left(H^{\prime}\right)\right|$, and the number of edges contained in graphs $H^{\prime \prime}$ and $H^{\prime} \backslash V\left(H^{\prime \prime}\right)$ is sufficiently large compared to $E\left(H^{\prime}\right)$. In the former two cases, we replace $H^{\prime}$ with graph $H^{\prime}\left[E^{\prime}\right]$ in $\mathcal{H}$, and graph $H^{\prime}\left[E^{\prime}\right]$ becomes inactive. In the latter case, we replace $H^{\prime}$ with two graphs: $H^{\prime \prime}$ and $H^{\prime} \backslash V\left(H^{\prime \prime}\right)$, that both remain active. The algorithm terminates once every graph in $\mathcal{H}$ is inactive. The crux of the analysis of the algorithm is to show that, when the algorithm terminates, the total number of edges lying in the subgraphs $H^{\prime} \in \mathcal{H}$ is high, compared to $|E(H)|$. This algorithm for decomposing graph $H$ into subgraphs and its analysis employ some of the techniques and ideas introduced in [CKN21, and is very similar in spirit to the hardness of approximation proof of the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning with Bundles problem, though details are different. We employ this decomposition algorithm multiple times, in order to obtain a partition $\left(E_{0}, E_{1}, \ldots, E_{z}\right)$ of the set $E(H)$ of edges into a small number of subsets, such that, among the constraints corresponding to the edges of $E_{0}$, at most a $1 / 4$-fraction can be satisfied by any assignment to $X \cup Y$, and, for all $1 \leq i \leq z$, a large fraction of constraints corresponding to edges of $E_{i}$ can be satisfied by some assignment. Depending on the cardinality of the set $E_{0}$ of edges we then determine whether $\mathcal{I}$ is a Yes-Instance or a No-Instance.

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

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

The main challenge is that, since the LP relaxation has $n^{\Theta(k)}$ variables, it is unclear how to solve it efficiently. We consider the dual linear program, that has poly $(n)$ variables and $n^{\Theta(k)}$ constraints. Using the $\alpha(n)$-approximation algorithm for Densest $k$-Subgraph as a subroutine, we design an approximate separation oracle for the dual LP, that allows us to solve the original LP-relaxation for Dense $k$-Coloring, obtaining a solution whose support size is bounded by poly $(n)$. By applying the LP-rounding approximation algorithm to this solution, we obtain the desired approximate solution to the input instance of Dense $k$-Coloring.

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

The two reductions are very similar; we focus on the reduction to Dense $k$-Coloring in this overview. Our construction is inspired by the results of KLS00, and we borrow some of our ideas from them. Assume that there is an efficient $\alpha(n)$-approximation algorithm for Dense $k$-Coloring. Let $G$ be an instance of the Densest $k$-Subgraph problem. The main difficulty in the reduction is that it is possible that $G$ only contains one very dense subgraph induced by $k$ vertices, while the Dense $k$-Coloring problem requires that the input graph $G$ can essentially be partitioned into many such dense subgraphs. To overcome this difficulty, we construct a random "inflated" bipartite graph $H$, that contains $n^{O(\log n)}$ vertices, where $n=|V(G)|$. Every vertex of $G$ is mapped to some vertex of $H$ at random, while every edge of $G$ is mapped to a large number of edges of $H$. This allows us to ensure that, if $G$ contains a subgraph $G^{\prime}$ induced by a set of $k$ vertices, where $\left|E\left(G^{\prime}\right)\right|=R$, then graph $H$ can essentially be partitioned into a large number of subgraphs that contain $k$ vertices each, and many of them contain close to $R$ edges. Therefore, we can apply our $\alpha(n)$-approximation algorithm for Dense $k$-Coloring to the new graph $H$. The main challenge in the reduction is that, while this approximation algorithm is guaranteed to return a large number of disjoint dense subgraphs of $H$, since every edge of $G$ contributes many copies to $H$, it is not clear that one can extract a single dense subgraph of $G$ from dense subgraphs of $H$. The main difficulty in the reduction is to ensure that, on the one hand, a single $k$-vertex dense subgraph in $G$ can be translated into $|V(H)| / k$ dense
subgraphs of $H$; and, on the other hand, a dense $k$-vertex subgraph of $H$ can be translated into a dense subgraph of $G$ on $k$ vertices. We build on and expand the ideas from KLS00 in order to ensure these properties.

Reductions between (r,h)-Graph Partitioning and Maximum Bounded-Crossing Subgraph. Lastly, we provide reductions between ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning and Maximum Bounded-Crossing Subgraph in both directions. First, we show that, if there is an efficient factor $\alpha(n)$-approximation algorithm for (r,h)-Graph Partitioning, then there is an efficient $O(\alpha(n) \cdot$ poly $\log n)$-approximation algorithm for Maximum Bounded-Crossing Subgraph. On the other hand, an efficient $\alpha(n)$-approximation algorithm for Maximum Bounded-Crossing Subgraph implies an efficient $O\left((\alpha(n))^{2} \cdot\right.$ poly $\left.\log n\right)$ approximation algorithm for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning. Combined with our conditional hardness of approximation for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, we get that, assuming the Low-Degree CSP Conjecture, for some constant $0<\varepsilon \leq 1 / 2$, there is no efficient $2^{(\log n)^{\varepsilon}}$-approximation algorithm for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, unless NP $\subseteq \operatorname{BPTIME}\left(n^{O(\log n)}\right)$.

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 $O(\sqrt{L+\Delta \cdot|E(G)|})$ edges, where $\Delta$ is maximum vertex degree in $G$. This result can be viewed as an extension of the classical Planar Separator Theorem of LT79. 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 $\mathcal{H}=\left\{H_{1}, \ldots, H_{r}\right\}$ 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 that show 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 one considers sufficiently large approximation factors, and allows randomized quasi-polynomial time algorithms). We also obtain conditional hardness of approximation results for all four problems based on the Low-Degree CSP Conjecture.

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 Section5 we provide reductions in the opposite direction. Lastly in Section 6 we provide reductions between (r,h)-Graph Partitioning and Maximum Bounded-Crossing Subgraph.

## 2 Preliminaries

By default, all logarithms are to the base of 2 . For a positive integer $N$, we denote by $[N]=$ $\{1,2, \ldots, N\}$. 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 $\operatorname{deg}_{G}(v)$ the degree of $v$ in $G$. For a subset $S$ of vertices of $G$, its volume is $\operatorname{vol}_{G}(S)=\sum_{v \in S} \operatorname{deg}_{G}(v)$. We sometimes omit the subscript $G$ if it is clear from the context.

Given a graph $G$, a drawing $\varphi$ 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 $e, e^{\prime}$ of $G$ cross at a point $p$, if $p \in \varphi(e) \cap \varphi\left(e^{\prime}\right)$, and $p$ is not the image of a shared endpoint of these edges. Given a graph $G$ and a drawing $\varphi$ of $G$ in the plane, we use $\operatorname{cr}(\varphi)$ to denote the number of crossings in $\varphi$, and the crossing number of $G$, denoted by $\operatorname{CrN}(G)$, is the minimum number of crossings in any drawing of $G$.

### 2.2 Problem Definitions and Additional Notation

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

Densest $k$-Subgraph. In the Densest $k$-Subgraph problem, the input is a graph $G$ and an integer $k>0$. The goal is to compute a subset $S \subseteq V(G)$ of $k$ vertices, maximizing $\left|E_{G}(S)\right|$. We denote an instance of the problem by $\operatorname{DkS}(G, k)$, and we denote the value of the optimal solution to instance $\operatorname{DkS}(G, k)$ by $\operatorname{OPT}_{\mathrm{DkS}}(G, k)$.
We also consider a bipartite version of the Densest $k$-Subgraph problem, called
Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph. This problem was first studied in $\left.\mathrm{AAM}^{+} 11\right]$. The input to the problem is a bipartite graph $G=(A, B, E)$ and positive integers $k_{1}, k_{2}$. The goal is to compute a subset $S \subseteq V(G)$ of vertices with $|S \cap A|=k_{1}$ and $|S \cap B|=k_{2}$, such that $\left|E_{G}(S)\right|$ is maximized. An instance of this problem is denoted by $\operatorname{BDkS}\left(G, k_{1}, k_{2}\right)$, and the value of the optimal solution to instance $\operatorname{BDkS}\left(G, k_{1}, k_{2}\right)$ is denoted by $\operatorname{OPT}_{\mathrm{BDkS}}\left(G, k_{1}, k_{2}\right)$. The following lemma shows that the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem and the Densest $k$-Subgraph problem are roughly equivalent from the approximation viewpoint. Similar results were also shown in prior work. For completeness, we provide the proof in Appendix A.

Lemma 2.1. Let $\alpha: \mathbb{Z}^{+} \rightarrow \mathbb{Z}^{+}$be an increasing function such that $\alpha(n)=o(n)$. Then the following hold:

- If there exists an $\alpha(n)$-approximation algorithm for the Densest $k$-Subgraph problem with running time at most $T(n)$, where $n$ is the number of vertices in the input graph, then there
exists an $O\left(\alpha\left(N^{2}\right)\right)$-approximation algorithm for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, with running time $O\left(T\left(N^{2}\right) \cdot \operatorname{poly}(N)\right)$, where $N$ is the number of vertices in the input graph. Moreover, if the algorithm for Densest $k$-Subgraph is deterministic, then so is the algorithm for Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph.
- Similarly, if there exists an efficient $\alpha(N)$-approximation algorithm for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, where $N$ is the number of vertices in the input graph, then there exists an efficient $O(\alpha(2 n))$-approximation algorithm for the Densest $k$-Subgraph problem, where $n$ is the number of vertices in the input graph. Moreover, if the algorithm for Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph is deterministic, then so is the algorithm for Densest $k$-Subgraph.

Dense $k$-Coloring. The input to the Dense $k$-Coloring problem consists of an $n$-vertex graph $G$ and an integer $k>0$, such that $n$ is an integral multiple of $k$. The goal is to compute a partition of $V(G)$ into $n / k$ subsets $S_{1}, \ldots, S_{n / k}$ of cardinality $k$ each, while maximizing $\sum_{i=1}^{n / k}\left|E_{G}\left(S_{i}\right)\right|$. An instance of the Dense $k$-Coloring problem is denoted by $\operatorname{DkC}(G, k)$, and the value of the optimal solution to instance $\operatorname{DkC}(G, k)$ is denoted by $\mathrm{OPT}_{\mathrm{DkC}}(G, k)$.
$(r, h)$-Graph Partitioning. The input to the (r,h)-Graph Partitioning problem consists of a graph $G$, and integers $r, h>0$. The goal is to compute $r$ vertex-disjoint subgraphs $H_{1}, \ldots, H_{r}$ of $G$, such that for all $1 \leq i \leq r,\left|E\left(H_{i}\right)\right| \leq h$, while maximizing $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right|$. An instance of the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning problem is denoted by $\operatorname{GP}(G, r, h)$, and the value of the optimal solution to instance $\operatorname{GP}(G, r, h)$ is denoted by $\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$.

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

### 2.3 Chernoff Bound

We use the following standard version of Chernoff Bound (see. e.g., [DP09).
Lemma 2.2 (Chernoff Bound). Let $X_{1}, \ldots, X_{n}$ be independent randon variables taking values in $\{0,1\}$. Let $X=\sum_{1 \leq i \leq n} X_{i}$, and let $\mu=\mathbf{E}[X]$. Then for any $t>2 e \mu$,

$$
\operatorname{Pr}[X>t] \leq 2^{-t}
$$

Additionally, for any $0 \leq \delta \leq 1$,

$$
\operatorname{Pr}[X<(1-\delta) \cdot \mu] \leq e^{-\frac{\delta^{2} \cdot \mu}{2}}
$$

### 2.4 Auxiliary Lemma

We use the following simple lemma.
Lemma 2.3. There is an efficient algorithm, that, given a graph $G$, a subset $S$ of its vertices, and a parameter $\frac{2}{|S|}<\beta<1$, computes a set $S^{\prime} \subseteq S$ of vertices, such that $\left|S^{\prime}\right| \leq \beta \cdot|S|$, and $\left|E_{G}\left(S^{\prime}\right)\right| \geq \Omega\left(\beta^{2} \cdot\left|E_{G}(S)\right|\right)$ holds.

Proof: Consider the graph $G^{\prime}=G[S]$ and denote $|S|=k$. We iteratively remove the lowestdegree vertex from $G^{\prime}$, until $G^{\prime}$ contains $\lfloor\beta k\rfloor$ vertices. Once the algorithm terminates, we output $S^{\prime}=V\left(G^{\prime}\right)$. It now remains to show that $\left|E_{G^{\prime}}\left(S^{\prime}\right)\right| \geq \Omega\left(\beta^{2}\left|E\left(G^{\prime}\right)\right|\right)$.
Observe that, if $H$ is an $n$-vertex graph, and $v$ is a lowest-degree vertex of $H$, then the degree of $v$ in $H$ is at most $2|E(H)| / n$. Therefore, if vertex $v$ is removed from $H$, then $|E(H)|$ decreases by at most a factor $(1-2 / n)$. Therefore,

$$
\frac{\left|E_{G^{\prime}}\left(S^{\prime}\right)\right|}{\left|E\left(G^{\prime}\right)\right|} \geq\left(1-\frac{2}{k}\right)\left(1-\frac{2}{k-1}\right) \cdots\left(1-\frac{2}{\lfloor\beta k\rfloor+1}\right)=\frac{\lfloor\beta k\rfloor \cdot(\lfloor\beta k\rfloor-1)}{k \cdot(k-1)}=\Omega\left(\beta^{2}\right) .
$$

## 3 Conditional Hardness of Densest $k$-Subgraph

### 3.1 Low-Degree CSP Conjecture

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

We associate with each constraint $C=C(x, y) \in \mathcal{C}$, a bipartite graph $G_{C}=(L, R, E)$, where $L=R=[A]$, and there is an edge $\left(a, a^{\prime}\right)$ in $E$ iff the assignments $a$ to $x$ and $a^{\prime}$ to $y$ satisfy $C$. Notice that instance $\mathcal{I}$ of the Bipartite 2-CSP problem is completely defined by $X, Y, A, \mathcal{C}$, and the graphs in $\left\{G_{C}\right\}_{C \in \mathcal{C}}$, so we will denote $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$. The size of instance $\mathcal{I}$ is defined to be size $(\mathcal{I})=|\mathcal{C}| \cdot A^{2}+|X|+|Y|$.
Consider some instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP. We say that $\mathcal{I}$ is a d-to-d instance if, for every constraint $C$, every vertex of graph $G_{C}=(L, R, E)$ has degree at most $d$.

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

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

### 3.2 Conditional Hardness of Densest $k$-Subgraph

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

Theorem 3.1. Assume that Conjecture 2holds and that $\mathrm{P} \neq \mathrm{NP}$. Then for some $0<\varepsilon \leq 1 / 2$, there is no efficient approximation algorithm for Densest $k$-Subgraph problem that achieves approximation factor $2^{(\log N)^{\varepsilon}}$, where $N$ is the number of vertices in the input graph.

In fact we will prove a slightly more general theorem, that will be useful for us later.
Theorem 3.2. Suppose there is an algorithm for the Densest $k$-Subgraph problem, that, given
 approximate solution to the problem, for some constant $0<\varepsilon \leq 1 / 2$. Then there is an algorithm, that, given an instance $\mathcal{I}$ of $(d(n), s(n))$ LD-2CSP problem of size $n$, where $d(n)=2^{(\log n)^{\varepsilon}}$ and $s(n)=1 / 2^{64(\log n)^{1 / 2+\varepsilon}}$, responds "YES" or "NO", in time $O(\operatorname{poly}(n) \cdot T(\operatorname{poly}(n)))$. If $\mathcal{I}$ is a YES-Instance, the algorithm is guaranteed to respond "YES", and if it is a No-Instance, it is guaranteed to respond "NO".

Theorem 3.1 immediately follows from Theorem 3.2. The remainder of this section is dedicated to proving Theorem 3.2. A central notion that we use is a constraint graph that is associated with an instance $\mathcal{I}$ of 2 -CSP.

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

Definition 3.3 (Bad Set of Constraints and Bad Collection of Edges). Let $\mathcal{C}^{\prime} \subseteq \mathcal{C}$ be a collection of constraints of an instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP. We say that $\mathcal{C}^{\prime}$ is a bad set of constraints if the largest number of constraints of $\mathcal{C}^{\prime}$ that can be simultaneously satisfied by
any assignment to the variables of $X \cup Y$ is at most $\frac{\left|\mathcal{C}^{\prime}\right|}{4}$. If $E^{\prime} \subseteq E(H(\mathcal{I}))$ is a set of edges of $H(\mathcal{I})$, whose corresponding set $\Phi\left(E^{\prime}\right)$ of constraints is bad, then we say that $E^{\prime}$ is a bad collection of edges.

The next observation easily follows from the definition of a bad set of constraints.
Observation 3.4. Let $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ be an instance of bipartite 2-CSP, and let $\mathcal{C}^{\prime}, \mathcal{C}^{\prime \prime} \subseteq$ $\mathcal{C}$ be two disjoint sets of constraints that are both bad. Then $\mathcal{C}^{\prime} \cup \mathcal{C}^{\prime \prime}$ is also a bad set of constraints.

Next, we define good subsets of constraints and good subgraphs of the constraint graph $H(\mathcal{I})$.
Definition 3.5 (Good Set of Constraints and Good Subgraphs of $H(\mathcal{I})$ ). Let $\mathcal{C}^{\prime} \subseteq \mathcal{C}$ be a collection of constraints of an instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP, and let $0<\beta \leq 1$ be a parameter. We say that $\mathcal{C}^{\prime}$ is a $\beta$-good set of constraints, if there is an assigmnet to variables of $X \cup Y$ that satisfies at least $\frac{\left|\mathcal{C}^{\prime}\right|}{\beta}$ constraints of $\mathcal{C}^{\prime}$ 。If $E^{\prime} \subseteq E(H(\mathcal{I}))$ is a set of edges of $H(\mathcal{I})$, whose corresponding set $\Phi\left(E^{\prime}\right)$ of constraints is $\beta$-good, then we say that $E^{\prime}$ is a $\beta$-good collection of edges. Lastly, if $H^{\prime} \subseteq H(\mathcal{I})$ is a subgraph of the constraint graph, and the set $E\left(H^{\prime}\right)$ of edges is $\beta$-good, then we say that graph $H^{\prime}$ is $\beta$-good.

The next observation easily follows from the definition of a good set of constraints.
Observation 3.6. Let $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ be an instance of bipartite 2-CSP, let $0<\beta \leq 1$ be a parameter, and let $H^{\prime}, H^{\prime \prime}$ be two subgraphs of $H(\mathcal{I})$ that are both $\beta$-good and disjoint in their vertices. Then graph $H^{\prime} \cup H^{\prime \prime}$ is also $\beta$-good.

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

The following theorem is key in proving Theorem 3.2.
Theorem 3.7. Assume that there exists a constant $0<\varepsilon \leq 1 / 2$, and an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the Densest $k$-Subgraph problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph, and $\alpha(N)=2^{(\log N)^{\varepsilon}}$. Then there is an algorithm, whose input consists of an instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP and parameter $n$ that is greater than a large enough constant, so that $\operatorname{size}(\mathcal{I}) \leq n$ holds, and $\mathcal{I}$ is a d(n)-to-d(n) instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^{\varepsilon}}$. Let $\beta=2^{8(\log n)^{1 / 2+\varepsilon}}$, and let $r=\lceil\beta \cdot \log n\rceil$. The algorithm returns a partition $\left(E^{b}, E_{1}, \ldots, E_{r}\right)$ of $E(H(\mathcal{I}))$, such that $E^{b}$ is a bad set of edges, and for all $1 \leq i \leq r$, set $E_{i}$ of edges is $\beta^{3}$-good. The running time of the algorithm is $O(T(\operatorname{poly}(n)) \cdot \operatorname{poly}(n)$.

The proof of Theorem 3.2 easily follows from Theorem 3.7. Assume that there exists a constant $0<\varepsilon \leq 1 / 2$, and an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the Densest $k$-Subgraph problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph, and $\alpha(N)=$ $2^{(\log N)^{\varepsilon}}$. We show an algorithm for the $(d(n), s(n))$-LD-2CSP problem, for $d(n)=2^{(\log n)^{\varepsilon}}$ and $s(n)=1 / 2^{64(\log n)^{1 / 2+\varepsilon}}$. Let $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ be an input instance of the Bipartite 2-CSP problem, with $\operatorname{size}(\mathcal{I})=n$, so that $\mathcal{I}$ is a $d(n)$-to- $d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^{\varepsilon}}$.

If $n$ is bounded by a constant, then we can determine whether $\mathcal{I}$ is a Yes-Instance or a NoInstance by exhaustively trying all assignments to its variables. Therefore, we assume that $n$ is greater than a large enough constant. We apply the algorithm from Theorem 3.7 to this instance $\mathcal{I}$. Let $\left(E^{b}, E_{1}, \ldots, E_{r}\right)$ be the partition of the edges of $E(H(\mathcal{I}))$ that the algorithm returns. We now consider two cases.

Assume first that $\left|E^{b}\right|>2|\mathcal{C}| / 3$. Let $\mathcal{C}^{b} \subseteq \mathcal{C}$ be the set of all constraints that correspond to the edges of $E^{b}$. Recall that set $\mathcal{C}^{b}$ of constraints is bad, so in any assignment, at most $\frac{\left|\mathcal{C}^{\prime}\right|}{4}$ of the constraints in $\mathcal{C}$ may be satisfied. Therefore, if $f$ is any assignment to variables of $X \cup Y$, the number of constraints in $\mathcal{C}$ that are not satisfied by $f$ is at least $\frac{3\left|\mathcal{C}^{\prime}\right|}{4}>\frac{|\mathcal{C}|}{2}$. Clearly, $\mathcal{I}$ may not be a Yes-Instance in this case. Therefore, if $\left|E^{b}\right|>2|\mathcal{C}| / 3$, we report that $\mathcal{I}$ is a No-Instance.
If $\left|E^{b}\right| \leq 2|\mathcal{C}| / 3$, then we report that $\mathcal{I}$ is a Yes-Instance. It is now enough to show that, if $\left|E^{b}\right| \leq 2|\mathcal{C}| / 3$, then instance $\mathcal{I}$ may not be a No-Instance. In other words, it is enough to show that there is an assignment that satisfies more than $\frac{|\mathcal{C}|}{2^{64(\log n)^{1 / 2+\varepsilon}}}$ constraints. Indeed, since $\left|E^{b}\right| \leq 2|\mathcal{C}| / 3$, there is an index $1 \leq i \leq r$, with $\left|E_{i}\right| \geq \frac{|\mathcal{C}|}{3 r}$. Since set $E_{i}$ of edges is $\beta^{3}$-good, there is an assignment to the variables of $X \cup Y$, that satisfies at least $\frac{\left|E_{i}\right|}{\beta^{3}} \geq \frac{|\mathcal{C}|}{3 r \beta^{3}}$ constraints that correspond to the edges of $E_{i}$. Recall that $\beta=2^{8(\log n)^{1 / 2+\varepsilon}}$ and $r=\lceil\beta \cdot \log n\rceil$. Therefore, $3 r \beta^{3} \leq 6 \beta^{4} \log n \leq 2^{64(\log n)^{1 / 2+\varepsilon}}$. We conclude that there is an assignment satisfying at least $|\mathcal{C}| / 2^{64(\log n)^{1 / 2+\varepsilon}}$ constraints, and so $\mathcal{I}$ may not be a No-Instance. It is easy to verify that the running time of the algorithm is $O(T(\operatorname{poly}(n)) \cdot \operatorname{poly}(n)$.

To conclude, we have shown that, if there is an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the Densest $k$-Subgraph problem, with running time at most $T(N)$, where $N$ is the number of vertices in the input graph, and $\alpha(N)=2^{(\log N)^{\varepsilon}}$, then there is an algorithm for the $(d(n), s(n))$-LD-2CSP problem, for $d(n)=2^{(\log n)^{\varepsilon}}$ and $s(n)=1 / 2^{8(\log n)^{1 / 2+\varepsilon}}$, whose running time is $O(T(\operatorname{poly}(n)) \cdot \operatorname{poly}(n)$.
In the remainder of this section we prove Theorem 3.7.

### 3.3 Proof of Theorem 3.7

The following theorem is the main technical ingredient of the proof of Theorem 3.7.
Theorem 3.8. Assume that there exists an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph. Then there is an algorithm, that, given an instance $\mathcal{I}=$ $\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP and parameters $n, \beta \geq 1$, so that $\operatorname{size}(\mathcal{I}) \leq n, \beta \geq$ $2^{30}(\alpha(n))^{3}(\log n)^{12}$, and $\mathcal{I}$ is a d(n)-to- $d(n)$ instance of Bipartite 2-CSP, for some function $d(n)$, in time $O(T(n) \cdot \operatorname{poly}(n))$, does one of the following:

- either correctly establishes that graph $H(\mathcal{I})$ is $\beta^{3}$-good; or
- computes a bad set $\mathcal{C}^{\prime} \subseteq \mathcal{C}$ of constraints, with $\left|\mathcal{C}^{\prime}\right| \geq \frac{|\mathcal{C}|}{8 \log ^{2} n}$; or
- computes a subgraph $H^{\prime}=\left(X^{\prime}, Y^{\prime}, E^{\prime}\right)$ of $H(\mathcal{I})$, for which the following hold:

$$
-\left|X^{\prime}\right| \leq \frac{2 d(n) \cdot|X|}{\beta} ;
$$

$-\left|Y^{\prime}\right| \leq \frac{2 d(n) \cdot|Y|}{\beta}$; and

$$
-\left|E^{\prime}\right| \geq \frac{\operatorname{vol}_{H}\left(X^{\prime} \cup Y^{\prime}\right)}{2048 d(n) \cdot \alpha(n) \cdot \log ^{4} n} .
$$

We prove Theorem 3.8 in Section Section 3.4, after we complete the proof of Theorem 3.7 using it. We start with the following corollary of Theorem 3.8.

Corollary 3.9. Assume that there exists an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the BDkS problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph. Then there is an algorithm, whose input consists of an instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP and parameters $n, \beta \geq 1$, so that $\operatorname{size}(\mathcal{I}) \leq n, \beta \geq 2^{30}(\alpha(n))^{3}(\log n)^{12}$, and $\mathcal{I}$ is a $d(n)-t o-d(n)$ instance of Bipartite 2-CSP. The algorithm returns a partition $\left(E_{1}, E_{2}\right)$ of $E(H(\mathcal{I}))$, where $E_{1}$ is a bad set of edges, and:

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

The running time of the algorithm is $O(T(n) \cdot \operatorname{poly}(n))$.
Proof: The algorithm is iterative. We start with $E_{1}=\emptyset, E_{2}=E(H(\mathcal{I}))$ and $H=H(\mathcal{I})$. We then iterate. In every iteration, we compute a graph $H^{\prime}=H \backslash E_{1}$. We denote by $\mathcal{C}^{\prime}=\Phi\left(E\left(H^{\prime}\right)\right)$ the set of all constraints of $\mathcal{C}$ corresponding to the edges of $H^{\prime}$. Notice that graph $H^{\prime}$ naturally defines a $d(n)$-to- $d(n)$ instance $\mathcal{I}^{\prime}$ of Bipartite 2-CSP, whose size is at most $n$, that corresponds to the subset $\mathcal{C}^{\prime} \subseteq \mathcal{C}$ of constraints. We apply the algorithm from Theorem 3.8 to instance $\mathcal{I}^{\prime}$. If the outcome of the algorithm is a bad set $\mathcal{C}^{\prime \prime} \subseteq \mathcal{C}^{\prime}$ of constraints, then we let $E=\left\{e_{C} \mid C \in \mathcal{C}^{\prime \prime}\right\}$ be the set of edges of $H^{\prime}$ corresponding to the constraints of $\mathcal{C}^{\prime \prime}$. We add the edges of $\tilde{E}$ to $E_{1}$, remove them from $E_{2}$, and continue to the next iteration.
If the algorithm from Theorem 3.8 certifies that graph $H^{\prime}$ is $\beta^{3}$-good, then we terminate the algorithm with the current partition $\left(E_{1}, E_{2}\right)$ of $E(H)$, and certify that the set $E_{2}$ of edges is $\beta^{3}$-good.

Otherwise, the outcome of the algorithm from Theorem 3.8 must be a subgraph $H^{\prime \prime}=\left(X^{\prime}, Y^{\prime}, E^{\prime}\right)$ of $H^{\prime}$, with $\left|X^{\prime}\right| \leq \frac{2 d(n) \cdot|X|}{\beta}$ and $\left|Y^{\prime}\right| \leq \frac{2 d(n) \cdot|Y|}{\beta}$. The algorithm also guarantees that $\left|E^{\prime}\right| \geq$ $\frac{\operatorname{vol}_{H^{\prime}}\left(X^{\prime} \cup Y^{\prime}\right)}{2048 d(n) \cdot \alpha(n) \cdot \log ^{4} n}$.
Let $E_{2}^{*}$ be the set of edges containing every edge $e \in E_{2}$ with exactly one endpoint in $V\left(H^{\prime \prime}\right)$. Since $E\left(H^{\prime}\right)=E_{2}$, it is immediate to verify that $\left|E_{2}^{*}\right| \leq \operatorname{vol}_{H^{\prime}}\left(X^{\prime} \cup Y^{\prime}\right)$. Therefore, we are guaranteed
that $\left|E^{\prime}\right| \geq \frac{\left|E_{2}^{*}\right|}{2048 d(n) \cdot \alpha(n) \cdot \log ^{4} n}$. We return the current partition $\left(E_{1}, E_{2}\right)$ of $E(H)$ and subgraph $H^{\prime \prime}$ of $H(\mathcal{I})$, and terminate the algorithm.
It is easy to verify that the algorithm consists of at most $O(\operatorname{poly}(n))$ iterations, and the running time of each iteration is at most $O(T(n) \cdot \operatorname{poly}(n))$. Therefore, the total running time of the algorithm is at most $O(T(n) \cdot \operatorname{poly}(n))$.
Next, we obtain the following corollary.
Corollary 3.10. Assume that there exists a constant $0<\varepsilon \leq 1 / 2$, and an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph, and $\alpha(N)=2^{(4 \log N)^{\varepsilon}}$. Then there is an algorithm, whose input consists of an instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP and parameter $n$ that is greater than a large enough constant, so that $\operatorname{size}(\mathcal{I}) \leq n$ holds, and $\mathcal{I}$ is a $d(n)-t o-d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^{\varepsilon}}$. Let $\beta=2^{8(\log n)^{1 / 2+\varepsilon}}$. The algorithm returns a partition $\left(E_{1}, E_{2}, E_{3}\right)$ of $E(H(\mathcal{I}))$, where $E_{1}$ is a bad set of constraints, $E_{2}$ is a $\beta^{3}$-good set of constraints, and $\left|E_{1} \cup E_{2}\right| \geq \frac{|E(H(\mathcal{I}))|}{\beta}$. The running time of the algorithm is $O(T(n) \cdot \operatorname{poly}(n))$.

Proof: Throughout the proof, we assume that there exists a constant $0<\varepsilon \leq 1 / 2$, and an $\alpha(N)$ approximation algorithm $\mathcal{A}$ for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph, and $\alpha(N)=2^{(2 \log N)^{\varepsilon}}$. Assume that we are given an instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP, together with a parameter $n$ that is greater than a large enough constant, so that $\operatorname{size}(\mathcal{I}) \leq n$, and $\mathcal{I}$ is a $d(n)$-to$d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^{\varepsilon}}$. For convenience, we denote $H=H(\mathcal{I})$. Our algorithm uses a parameter $\eta=2^{12} d(n) \cdot \alpha(n) \cdot \log ^{4} n$.
The algorithm is iterative. Over the course of the algorithm, we maintain a collection $\mathcal{H}$ of subgraphs of $H$, and another subgraph $H^{g}$ of $H$. We will ensure that, throughout the algorithm, all graphs in $\mathcal{H} \cup\left\{H^{g}\right\}$ are mutually disjoint in their vertices. We denote by $E^{g}=E\left(H^{g}\right)$ and $E^{1}=\bigcup_{H^{\prime} \in \mathcal{H}} E\left(H^{\prime}\right)$. Additionally, we maintain another set $E^{b}$ of edges of $H$, that is disjoint from $E^{g} \cup E^{1}$, and we denote by $E^{0}=E(H) \backslash\left(E^{g} \cup E^{b} \cup E^{1}\right)$ the set of all remaining edges of $H$. We ensure that the following invariants hold throughout the algorithm.

I1. set $E^{g}=E\left(H^{g}\right)$ of edges is $\beta^{3}$-good;
I2. set $E^{b}$ of edges is bad; and
I3. all graphs in $\mathcal{H} \cup\left\{H^{g}\right\}$ are disjoint in their vertices.
Intuitively, we will start with the set $\mathcal{H}$ containing a single graph $H$, and $E^{g}=E^{b}=E^{0}=\emptyset$. As the algorithm progresses, we will iteratively add edges to sets $E^{g}, E^{b}$ and $E^{0}$, while partitioning the graphs in $\mathcal{H}$ into smaller subgraphs. The algorithm will terminate once $\mathcal{H}=\emptyset$. The key in the analysis of the algorithm is to ensure that $\left|E^{0}\right|$ is relatively small when the algorithm terminates. We do so via a charging scheme: we assign a budget to every edge of $E^{1} \cup E^{g} \cup E^{b}$, that evolves over the course of the algorithm, and we keep track of this budget over the course of the algorithm.
In order to define vertex budgets, we will assign, to every graph $H \in \mathcal{H}$ a level, that is an integer between 0 and $\lceil\log n\rceil$. We will ensure that, throughout the algorithm, the following additional invariants hold:

I4. If $H^{\prime} \in \mathcal{H}$ is a level- $i$ graph, then the budget of every edge $e \in E\left(H^{\prime}\right)$ is at most $\eta^{i}$; and
I5. Throughout the algorithm's execution, the total budget of all edges in $E^{g} \cup E^{b} \cup E^{1}$ is at least $|E(H)|$.

Intuitively, at the end of the algorithm, we will argue that the level of every graph in $\mathcal{H}$ 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 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 $\left|E^{g} \cup E^{b} \cup E^{1}\right|$ is sufficiently large. We now proceed to describe the algorithm.

Our algorithm will repeatedly use the algorithm from Corollary 3.9 , with the same functions $\alpha(N), d(n)$, and parameter $\beta$. In order to be able to use the corollary, we need to estalish that $\beta \geq 2^{30}(\alpha(n))^{3}(\log n)^{12}$. This is immediate to verify since $\beta=2^{8(\log n)^{1 / 2+\varepsilon}}, \alpha(n)=2^{(4 \log n)^{\varepsilon}}$, and $n$ is large enough.

Initialization. At the beginning of the algorithm, we set $E^{0}=E^{g}=E^{b}=\emptyset$, and we let $\mathcal{H}$ contain a single graph $H$, which is assigned level 0 . Note that $E^{1}=E(H)$ must hold. Every edge $e \in E(H)$ is assigned budget $b(e)=1$. Clearly, the total budget of all edges of $E^{1} \cup E^{g} \cup E^{b}$ is $B=\sum_{e \in E^{1} \cup E^{g} \cup E^{b}} b(e)=|E(H)|$.

The algorithm performs iterations, as long as $\mathcal{H} \neq \emptyset$. In every iteration, we select an arbitrary graph $H^{\prime} \in \mathcal{H}$ to process. We now describe a single iteration.

Iteration description. We now describe an iteration where some graph $H^{\prime} \in \mathcal{H}$ is processed. We assume that graph $H^{\prime}$ is assigned level $i$. Notice that graph $H^{\prime}$ naturally defines an instance $\mathcal{I}^{\prime}=\left(X^{\prime}, Y^{\prime}, A, \mathcal{C}^{\prime},\left\{G_{C}\right\}_{C \in \mathcal{C}^{\prime}}\right)$ of Bipartite 2-CSP, where $X^{\prime}=V\left(H^{\prime}\right) \cap X, Y^{\prime}=V\left(H^{\prime}\right) \cap Y$, $\mathcal{C}^{\prime}=\left\{C \in \mathcal{C} \mid e_{C} \in E\left(H^{\prime}\right)\right\}$, and the graphs $G_{C}$ for constraints $C \in \mathcal{C}^{\prime}$ remain the same as in instance $\mathcal{I}$. Clearly, $\operatorname{size}\left(\mathcal{I}^{\prime}\right) \leq \operatorname{size}(\mathcal{I}) \leq n$, and $H\left(\mathcal{I}^{\prime}\right)=H^{\prime}$. Furthermore, instance $\mathcal{I}^{\prime}$ remains a $d(n)$-to- $d(n)$ instance. We apply the algorithm from Corollary 3.9 to instance $\mathcal{I}^{\prime}$, with parameters $n$ and $\beta$ remaining unchanged. Consider the partition $\left(E_{1}, E_{2}\right)$ of $E\left(H^{\prime}\right)$ that the algorithm returns. Recall that the set $E_{1}$ of edges is bad. We add the edges of $E_{1}$ to set $E^{b}$. From Invariant I2 and Observation 3.4, set $E^{b}$ of edges continues to be bad. If the algorithm from Corollary 3.9 certified that $E_{2}$ is a $\beta^{3}$-good set of edges, then we update graph $H^{g}$ to be $H^{g} \cup\left(H^{\prime} \backslash E_{1}\right)$, and we add the edges of $E_{2}$ to set $E^{g}$. We then remove graph $H^{\prime}$ from $\mathcal{H}$, and continue to the next iteration. Note that, from Observation 3.6 and Invariants $I 1$ and $I 3$, the set $E^{g}$ of edges continues to be $\beta^{3}$-good. It is easy to verify that all remaining invariants also continue to hold.

From now on we assume that the algorithm from Corollary 3.9 returned a subgraph $H^{\prime \prime}=\left(X^{\prime \prime}, Y^{\prime \prime}, E^{\prime \prime}\right)$ of $H^{\prime}$, with $E^{\prime \prime} \subseteq E_{2}$, such that $\left|X^{\prime \prime}\right| \leq \frac{2 d(n) \cdot\left|X^{\prime}\right|}{\beta}$ and $\left|Y^{\prime \prime}\right| \leq \frac{2 d(n) \cdot\left|Y^{\prime}\right|}{\beta}$. In particular, $\left|V\left(H^{\prime \prime}\right)\right|=$ $\left|X^{\prime \prime}\right|+\left|Y^{\prime \prime}\right| \leq \frac{2 d(n)}{\beta} \cdot\left(\left|X^{\prime}\right|+\left|Y^{\prime}\right|\right) \leq \frac{2 d(n)}{\beta} \cdot\left|V\left(H^{\prime}\right)\right|$. Additinally, if we denote by $E_{2}^{*}$ the subset of edges of $E_{2}$ containing all edges with exactly one endpoint in $X^{\prime \prime} \cup Y^{\prime \prime}$, then $\left|E^{\prime \prime}\right| \geq \frac{\left|E_{2}^{*}\right|}{2048 d(n) \cdot \alpha(n) \cdot \log ^{4} n}$ must hold. We let $H^{*}$ be the graph obtained from $H^{\prime} \backslash E_{1}$, by deleting the vertices of $H^{\prime \prime}$ from it, so $V\left(H^{*}\right) \cup V\left(H^{\prime \prime}\right)=V\left(H^{\prime}\right)$, and $E\left(H^{*}\right) \cup E\left(H^{\prime \prime}\right) \cup E_{2}^{*}=E_{2}$. We remove graph $H^{\prime}$ from $\mathcal{H}$, and we add graphs $H^{\prime \prime}$ and $H^{*}$ to $\mathcal{H}$, with graph $H^{\prime \prime}$ assigned level $(i+1)$, and graph $H^{*}$ assigned level $i$. We also add the edges of $E_{2}^{*}$ to $E^{0}$, and we update the set $E^{1}$ of edges to contain all edges of
$\bigcup_{\tilde{H} \in \mathcal{H}} E(\tilde{H})$. Since we did not modify graph $H^{g}$ in the current iteration, it is immediate to verify that Invariants I1 $I 3$ continue to hold. Next, we update the budgets of edges, in order to ensure that Invariants I4 and I5 continue to hold. Intuitively, the edges of $E_{2}^{*}$ are now added to set $E^{0}$, so we need to distribute their budget among the edges of $E\left(H^{\prime \prime}\right)$, in order to ensure that the total budget of all edges in $E^{g} \cup E^{b} \cup E^{1}$ does not decrease. This will ensure that Invariant 15 continues to hold. At the same time, since the level of graph $H^{\prime \prime}$ is $(i+1)$, while the level of graph $H^{\prime}$ was $i$, we can increase the budgets of the edges of $E\left(H^{\prime}\right)$ and still maintain Invariant I4
Formally, recall that Corollary 3.9 guarantees that $\left|E_{2}^{*}\right| \leq\left|E^{\prime \prime}\right| \cdot\left(2048 d(n) \cdot \alpha(n) \cdot \log ^{4} n\right)=\frac{\left|E^{\prime \prime}\right| \cdot \eta}{2}$. From Invariant I4, the current budget of every edge in $E^{\prime \prime} \cup E_{2}^{*}$ is bounded by $\eta^{i}$. Therefore, at the beginning of the current iteration:

$$
\sum_{e \in E^{\prime \prime} \cup E_{2}^{*}} b(e) \leq \eta^{i} \cdot\left(\left|E_{2}^{*}\right|+\left|E^{\prime \prime}\right|\right) \leq \eta^{i} \cdot\left|E^{\prime \prime}\right| \cdot\left(1+\frac{\eta}{2}\right)<\eta^{i+1} \cdot\left|E^{\prime \prime}\right| .
$$

We set the budget of every edge in $E^{\prime \prime}$ to be $\eta^{i+1}$, and leave the budgets of all other 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 iteration, so Invariant I5 continues to hold. It is also easy to verify that Invariant $[4$ continues to hold. Therefore, all invariants continue to hold at the end of the iteration. This completes the description of an iteration.
The algorithm terminates when $\mathcal{H}=\emptyset$. Clearly, we obtain a partition $\left(E^{g}, E^{b}, E^{0}\right)$ of $E(H)$ into disjoint subsets, where the set $E^{b}$ of edges is bad, and the set $E^{g}$ of edges is $\beta^{3}$-good. It remains to show that $\left|E^{g} \cup E^{b}\right| \geq \frac{|E(H)|}{\beta}$. We use the edge budgets in order to prove this. Let $L^{*}$ be the largest level of any subgraph of $H$ that belonged to $\mathcal{H}$ at any time during the algorithm. We start with the following key observation.

Observation 3.11. $L^{*} \leq(\log n)^{1 / 2-\varepsilon}$.

Proof: Consider any graph $H^{\prime \prime}$ that was added to set $\mathcal{H}$ at any time during the algorithm's execution, and assume that $H^{\prime \prime}$ was assigned level $i$. Consider the iteration during which $H^{\prime \prime}$ was added to $\mathcal{H}$, and let $H^{\prime} \in \mathcal{H}$ be the graph that was processed during that iteration. We refer to graph $H^{\prime}$ as the parent-graph of $H^{\prime \prime}$. Note that the level of $H^{\prime}$ is either $i$ or $(i-1)$. Assume that it is the latter. Then, from the algorithm's description, $\left|V\left(H^{\prime \prime}\right)\right| \leq \frac{2 d(n)}{\beta} \cdot\left|V\left(H^{\prime}\right)\right|$ must hold.

We can now construct a partitioning tree, that contains a vertex $v\left(H^{\prime}\right)$ for every graph $H^{\prime}$ that was ever present in $\mathcal{H}$ over the course of the algorithm, an an edge between vertices $v\left(H^{\prime}\right)$ and $v\left(H^{\prime \prime}\right)$ whenever graph $H^{\prime}$ is a parent-graph of graph $H^{\prime \prime}$. The root of the tree is $v(H)$. Consider now again some graph $H^{\prime \prime}$, and the unique path $P$ in the partitioning tree, connecting $v(H)$ to $v\left(H^{\prime \prime}\right)$. Denote the vertices on this path by $v(H)=v\left(H_{0}\right), v\left(H_{1}\right), \ldots, v\left(H_{r}\right)=v\left(H^{\prime \prime}\right)$, and assume that these vertices appear on path $P$ in this order. For all $1 \leq i \leq r$, denote the level of graph $H_{i}$ by $L_{i}$. Then $0=L_{1} \leq L_{2} \leq \cdots \leq L_{r}$ must hold. Moreover, for every index $0<i \leq r$, either $L_{i}=L_{i-1}$; or $L_{i}=L_{i-1}+1$ hold. In the latter case, $\left|V\left(H_{i}\right)\right| \leq\left|V\left(H_{i-1}\right)\right| \cdot \frac{2 d(n)}{\beta}$. Denote $\Delta=\frac{\log n}{\log \left(\frac{\beta}{2 d(n)}\right)}$.
We claim that $L_{r} \leq \Delta$. Indeed, assume for contradiction that $L_{r}>\Delta$. Then there is a collection $J \subseteq\{1, \ldots, r\}$ of at least $\Delta+1$ indices $i$, for which $L_{i}=L_{i-1}+1$. But then:

$$
\left|V\left(H^{\prime \prime}\right)\right| \leq n \cdot\left(\frac{2 d(n)}{\beta}\right)^{\Delta+1}<1
$$

a contradiction. We conclude that $L^{*} \leq \frac{\log n}{\log \left(\frac{\beta}{2 d(n)}\right)}$. Substituting $\beta=2^{8(\log n)^{1 / 2+\varepsilon}}$ and $d(n) \leq$ $2^{(\log n)^{\varepsilon}}$, we get that:

$$
L^{*} \leq \frac{\log n}{\log \left(2^{6(\log n)^{1 / 2+\varepsilon}}\right)} \leq(\log n)^{1 / 2-\varepsilon} .
$$

From Invariant I4 throughout the algorithm, for every edge $e \in E^{1}, b(e) \leq \eta^{L^{*}}$ must hold. Once an edge is added to $E^{b} \cup E^{g}$, its budget does not change. Therefore, at the end of the algorithm, the budget of every edge in $E^{g} \cup E^{b}$ is at most $\eta^{L^{*}}$. On the other hand, from Invariant $I 5$, at the end of the algorithm, the total budget of all edges in $E^{1} \cup E^{g} \cup E^{b}$ is at least $|E(H)|$. Therefore, at the end of the algorithm:

$$
\left|E^{g} \cup E^{b}\right| \geq \frac{|E(H)|}{\eta^{L^{*}}}
$$

We now bound $\eta^{L^{*}}$. Recall that $\eta=2^{12} d(n) \cdot \alpha(n) \cdot \log ^{4} n \leq 2^{4(\log n)^{\varepsilon}}$, since $d(n) \leq 2^{(\log n)^{\varepsilon}}$, $\alpha(n)=2^{(4 \log n)^{\varepsilon}}$, and $n$ is large enough. Since, from Observation 3.11, $L^{*} \leq(\log n)^{1 / 2-\varepsilon}$, we get that $\eta^{L^{*}} \leq 2^{4(\log n)^{1 / 2}}<\beta$, since $\beta=2^{8(\log n)^{1 / 2+\varepsilon}}$. Therefore, $\left|E^{g} \cup E^{b}\right| \geq|E(H)| / \beta$ as required.

Lastly, it is easy to verify that the algorithm has at most poly $(n)$ iterations, and the running time of each iteration is bounded by $O(T(n) \cdot \operatorname{poly}(n))$, so the total running time of the algorithm is at most $O(T(n) \cdot \operatorname{poly}(n))$.
We are now ready to complete the proof of Theorem 3.7. Assume that there exists a constant $0<\varepsilon \leq 1 / 2$, and an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the Densest $k$-Subgraph problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph, and $\alpha(N)=$ $2^{(\log N)^{\varepsilon}}$. From Lemma 2.1, there exists an $\alpha^{\prime}(N)$-approximation algorithm $\mathcal{A}$ for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, where $N$ is the number of vertices in the input graph, and $\alpha^{\prime}(N) \leq O\left(\alpha\left(N^{2}\right)\right) \leq O\left(2^{(2 \log N)^{\varepsilon}}\right)$. The running time of the algorithm is at most $O\left(T\left(N^{2}\right)\right.$. $\operatorname{poly}(N))$. Denote $T^{\prime}(N)=O\left(T\left(N^{2}\right) \cdot \operatorname{poly}(N)\right)$ this bound on the running time of the algorithm, and let $\alpha^{\prime \prime}(N)=2^{(4 \log N)^{\varepsilon}}$. Then there is an $\alpha^{\prime \prime}(N)$-approximation algorithm for Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph with running time at most $O\left(T^{\prime}(N)\right)$. Indeed, if $N$ is greater than a sufficiently large constant, then we can use Algorithm $\mathcal{A}$, to obtain a solution whose approximation factor is $\alpha^{\prime}(N) \leq O\left(2^{2 \log N)^{\varepsilon}}\right) \leq 2^{(4 \log N)^{\varepsilon}} \leq \alpha^{\prime \prime}(N)$. Otherwise, we can solve the problem exactly via exhaustive search.

Assume now that we are given an instance $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ of Bipartite 2-CSP and parameter $n$ that is greater than a large enough constant, so that $\operatorname{size}(\mathcal{I}) \leq n$ holds, and $\mathcal{I}$ is a $d(n)$-to- $d(n)$ instance of Bipartite 2-CSP, for $d(n) \leq 2^{(\log n)^{\varepsilon}}$. Let $\beta=2^{8(\log n)^{1 / 2+\varepsilon}}$, and let $r=\lceil\beta \cdot \log n\rceil$. For convenience, we denote $H=H(\mathcal{I})$. Initially, we set $E^{b}=\emptyset$. Our algorithm performs $r$ iterations, where for all $1 \leq j \leq r$, in iteration $j$ we construct the set $E_{j} \subseteq E(H)$ of
edges, that is $\beta^{3}$-good, and possibly adds some edges to set $E^{b}$. We ensure that, throughout the algorithm, the set $E^{b}$ of edges is bad.
Initially, $E^{b}=\emptyset$. We now describe the $j$ th iteration. We assume that sets $E_{1}, \ldots, E_{j-1}$ of edges of $H$ were already defined. We construct graph $H_{j}$, that is obtained from graph $H$, by deleting the edges of $E_{1} \cup \cdots \cup E_{j-1} \cup E^{b}$ from it. Notice that graph $H_{j}$ naturally defines an instance $\mathcal{I}_{j}=$ $\left(X, Y, A, \mathcal{C}_{j},\left\{G_{C}\right\}_{C \in \mathcal{C}_{j}}\right)$ of Bipartite 2-CSP, with $H_{j}=H\left(\mathcal{I}_{j}\right)$, where $\mathcal{C}_{j}=\left\{C \in \mathcal{C} \mid e_{C} \in E\left(H_{j}\right)\right\}$. We apply the algorithm from Corollary 3.10 to graph $H_{j}$, with parameters $n, \beta$, and $d(n)$ remaining unchanged. Consider a partition $\left(E^{1}, E^{2}, E^{3}\right)$ of $E\left(H_{j}\right)$ that the algorithm returns. We add the edges of $E^{1}$ to set $E^{b}$. Since both sets of edges are bad, from Observation 3.4, set $E^{b}$ of edges continues to be bad. We also set $E_{j}=E^{2}$, which is guaranteed to be a $\beta^{3}$-good set of edges from Corollary 3.10. Recall that Corollary 3.10 also guarantees that $\left|E^{1} \cup E^{2}\right| \geq\left|E\left(H_{j}\right)\right| / \beta$. We then continue to the next iteration.
Since, from the above discussion, for all $1 \leq j<r,\left|E\left(H_{j+1}\right)\right| \leq\left(1-\frac{1}{\beta}\right)\left|E\left(H_{j}\right)\right|$, and since $r=\lceil\beta \cdot \log n\rceil$, at the end of the algorithm, we are guaranteed that the final collection $E^{b}, E_{1}, \ldots, E_{r}$ of subsets of edges indeed partitions $E(H)$.
Notice that the running time of a single iteration is bounded by $O\left(T^{\prime}(n) \cdot \operatorname{poly}(n)\right) \leq O(T(\operatorname{poly}(n))$. $\operatorname{poly}(n))$. Since the number of iterations is bounded by poly $(n)$, the total running time of the algorithm is bounded by $O(T(\operatorname{poly}(n)) \cdot \operatorname{poly}(n))$.
In order to complete the proof of Theorem 3.7, it is now enough to prove Theorem 3.8, which we do next.

### 3.4 Proof of Theorem 3.8

The proof partially relies on ideas and techniques from [CKN21. Assume that there exists an $\alpha(N)$-approximation algorithm $\mathcal{A}$ for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, whose running time is at most $T(N)$, where $N$ is the number of vertices in the input graph. Let $\mathcal{I}=\left(X, Y, A, \mathcal{C},\left\{G_{C}\right\}_{C \in \mathcal{C}}\right)$ be the input instance of Bipartite 2-CSP, with $\operatorname{size}(\mathcal{I}) \leq n$. For convenience, we denote $H=H(\mathcal{I})$. If $|E(H)| \leq \beta^{3}$, then graph $H$ is $\beta^{3}$-good, since we can compute an assignment to the variables of $X \cup Y$ that satisfies at least one constraint of $\mathcal{C}$. Therefore, we assume from now on that $|E(H)|>\beta^{3}$. We can also assume that graph $H$ contains no isolated vertices, as isolated vertices of $H$ correspond to variables that do not participate in any constraints, and can be discarded.

The proof consists of four steps. In the first step, in order to simplify the proof, we will regularize graph $H$, by computing a "nice" subgraph $\tilde{H} \subseteq H$. In the second step, we will define an assignment graph associated with $\tilde{H}$, and we will use it in order to obtain an instance of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph, to which algorithm $\mathcal{A}$ will then be applied. In the last two steps, we will use the outcome of algorithm $\mathcal{A}$ in order to either correctly establish that graph $H$ is $\beta^{3}$-good, or to compute a bad subset of constraints, or a subgraph $H^{\prime}$ of $H$ as required. We now describe each of the three steps in turn.

### 3.4.1 Step 1: Regularization

In this step we will compute a subgraph $\tilde{H}$ of $H$ that has a convenient structure. We refer to graphs with such structure as nice subgraphs of $H$, and define them next.

Definition 3.12 (Nice Subgraph of $H$ ). Let $\tilde{H}=(\tilde{X}, \tilde{Y}, \tilde{E})$ be a subgraph of $H$, and let $d_{1}, d_{2} \geq 1$ be parameters. We say that $\tilde{H}$ is a $\left(d_{1}, d_{2}\right)$-nice subgraph of $H$, if the following hold:

- For every vertex $x \in \tilde{X}, d_{1} \leq \operatorname{deg}_{H}(x)<2 d_{1}$;
- For every vertex $y \in \tilde{Y}, d_{2} \leq \operatorname{deg}_{H}(y)<8 d_{2} \log n$ and $d_{2} \leq \operatorname{deg}_{\tilde{H}}(y)<2 d_{2}$; and
- $|\tilde{E}| \geq \frac{d_{1}}{4 \log n} \cdot|\tilde{X}|$.

We say that $\tilde{H}$ is a nice subgraph of $H^{\prime}$ if it is a $\left(d_{1}, d_{2}\right)$-nice subgraph of $H$ for any pair $d_{1}, d_{2} \geq 1$ of parameters.

The first step of our algorithm is summarized in the following claim, that allows us to compute a nice subgraph $\tilde{H}$ of $H$ that contains many edges of $H$.

Claim 3.13. There is an algorithm with running time $O(\operatorname{poly}(n))$, that computes parameters $d_{1}, d_{2}>1$, and a subgraph $\tilde{H}$ of $H$, such that $\tilde{H}$ is a $\left(d_{1}, d_{2}\right)$-nice subgraph of $H$, and $|E(\tilde{H})| \geq$ $\frac{|E(H)|}{8 \log ^{2} n}$.

Proof: The proof uses standard regularization techniques, and consists of three steps. Denote $H=(X, Y, E)$.

In the first step, we partition the vertices of $X$ into groups $S_{0}, \ldots, S_{q}$, for $q=\lceil\log n\rceil$, where for all $0 \leq i \leq q, S_{i}=\left\{x \in X \mid 2^{i} \leq \operatorname{deg}_{H}(x)<2^{i+1}\right\}$. We also partition the set $E$ of edges into subsets $E_{0}, \ldots, E_{q}$, where for all $0 \leq i \leq q$, set $E_{i}$ contains all edges $e \in E$ that are incident to vertices of $S_{i}$. Clearly, there is an index $0 \leq i^{*} \leq q$, with $\left|E_{i^{*}}\right| \geq \frac{|E(H)|}{2 \log n}$. We let $\tilde{X}=S_{i^{*}}$, and we let $H_{1}$ be the graph whose vertex set is $\tilde{X} \cup Y$, and edge set is $E_{i^{*}}$. We also define $d_{1}=2^{i^{*}}$. Clearly, for every vertex $x \in \tilde{X}, d_{1} \leq \operatorname{deg}_{H}(x)<2 d_{1}$. This completes the first regularization step.
We now proceed to describe our second step, in which we consider the vertices of $y \in Y$ one by one. We say that a vertex $y \in Y$ is bad, if $\operatorname{deg}_{H_{1}}(y)<\frac{\operatorname{deg}_{H}(y)}{4 \log n}$. Let $Y^{\prime \prime} \subseteq Y$ be the set of all bad vertices, and let $Y^{\prime}=Y \backslash Y^{\prime \prime}$ be the set of all remaining vertices of $Y$, that we refer to as good vertices. We use the following observation.
Observation 3.14. $\sum_{y \in Y^{\prime \prime}} \operatorname{deg}_{H_{1}}(y) \leq \frac{\left|E\left(H_{1}\right)\right|}{2}$.
Proof: Since, for every bad vertex $y, \operatorname{deg}_{H_{1}}(y)<\frac{\operatorname{deg}_{H}(y)}{4 \log n}$, we get that:

$$
\sum_{y \in Y^{\prime \prime}} \operatorname{deg}_{H_{1}}(y)<\sum_{y \in Y^{\prime \prime}} \frac{\operatorname{deg}_{H}(y)}{4 \log n} \leq \frac{|E(H)|}{4 \log n}
$$

Since, as observed above, $\left|E\left(H_{1}\right)\right| \geq \frac{|E(H)|}{2 \log n}$, the observation follows.

We let $H_{2}$ be a graph that is obtained from $H_{1}$, by discarding the vertices of $Y^{\prime \prime}$ from it. Therefore, $V\left(H_{2}\right)=\tilde{X} \cup Y^{\prime}$. Additionally, from Observation 3.14. $\left|E\left(H_{2}\right)\right| \geq \frac{\left|E\left(H_{1}\right)\right|}{2} \geq \frac{|E(H)|}{4 \log n}$.
Lastly, in our third step, we perform a geometric grouping of the vertices of $Y^{\prime}$ by their degree in $H_{2}$. Specifically, we let $r=\lceil\log n\rceil$, and we partition the vertices of $Y^{\prime}$ into sets $S_{0}^{\prime}, \ldots, S_{r}^{\prime}$, where for $0 \leq j \leq r, S_{j}^{\prime}=\left\{y \in Y^{\prime} \mid 2^{j} \leq \operatorname{deg}_{H_{2}}(y)<2^{j+1}\right\}$. As before, we also partition the set $E\left(H_{2}\right)$ of edges into subsets $E_{0}^{\prime}, \ldots, E_{r}^{\prime}$, where for $0 \leq j \leq r$ set $E_{j}^{\prime}$ contains all edges $e \in E\left(H_{2}\right)$ that are incident to the vertices of $S_{j}^{\prime}$. As before, there must be an index $0 \leq j^{*} \leq r$ with $\left|E_{j^{*}}^{\prime}\right| \geq \frac{\left|E\left(H_{2}\right)\right|}{2 \log n} \geq \frac{|E(H)|}{8 \log ^{2} n}$. We set $\tilde{Y}=S_{j^{*}}^{\prime}, d_{2}=2^{j^{*}}$, and we let $\tilde{H}$ be the graph whose vertex set is $\tilde{X} \cup \tilde{Y}$, and edge set is $E_{j^{*}}^{\prime}$. We now verify that this graph has all required properties.
First, as observed already, for every vertex $x \in \tilde{X}, d_{1} \leq \operatorname{deg}_{H}(x)<2 d_{1}$. Let $\tilde{E}=E(\tilde{H})$. As observed already, $|\tilde{E}| \geq \frac{|E(H)|}{8 \log ^{2} n}$. Moreover, since, for every vertex $x \in \tilde{X}, \operatorname{deg}_{H_{1}}(x)=\operatorname{deg}_{H}(x) \geq d_{1}$, we get that $\left|E\left(H_{1}\right)\right| \geq d_{1} \cdot|\tilde{X}|$, and so $|\tilde{E}| \geq \frac{\left|E\left(H_{2}\right)\right|}{2 \log n} \geq \frac{\left|E\left(H_{1}\right)\right|}{4 \log n} \geq \frac{d_{1}}{4 \log n} \cdot|\tilde{X}|$.
Consider now some vertex $y \in \tilde{Y}$. From the definition of graph $\tilde{H}$, it is immediate to verify that $\operatorname{deg}_{\tilde{H}}(y)=\operatorname{deg}_{H_{2}}(y)$. Therefore, $d_{2} \leq \operatorname{deg}_{\tilde{H}}(y)<2 d_{2}$. Clearly, $\operatorname{deg}_{H}(y) \geq \operatorname{deg}_{\tilde{H}}(y) \geq d_{2}$. Lastly, since vertex $y$ is good, we get that:

$$
\operatorname{deg}_{\tilde{H}}(y)=\operatorname{deg}_{H_{2}}(y)=\operatorname{deg}_{H_{1}}(y) \geq \frac{\operatorname{deg}_{H}(y)}{4 \log n} .
$$

Since $\operatorname{deg}_{\tilde{H}}(y)<2 d_{2}$, we get that $\operatorname{deg}_{H}(y) \leq(4 \log n) \operatorname{deg}_{\tilde{H}}(y)<8 d_{2} \log n$. We conclude that $d_{2} \leq \operatorname{deg}_{H}(y)<8 d_{2} \log n$.

### 3.4.2 Step 2: Assignment Graph and Reduction to Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph

Recall that we have computed, in the first step, a subgraph $\tilde{H}=(\tilde{X}, \tilde{Y}, \tilde{E})$ of the graph $H=\underset{\sim}{H}(\mathcal{I})$. Since every edge of $H$ is associated with a distinct constraint in $\mathcal{C}$, we can define a collection $\tilde{\mathcal{C}} \subseteq \mathcal{C}$ of constraints corresponding to the edges of $\tilde{H}: \tilde{\mathcal{C}}=\left\{C \in \mathcal{C} \mid e_{C} \in \tilde{E}\right\}$.
Next, we define a bipartite graph $G=(U, V, \hat{E})$, called assignment graph, that is associated with graph $\tilde{H}$. For every variable $z \in \tilde{X} \cup \tilde{Y}$, we define a set $R(z)=\{v(z, a) \mid 1 \leq a \leq A\}$ of vertices that represent the possible assignments to variable $z$. We then set $U=\bigcup_{x \in \tilde{X}} R(x)$, and $V=\bigcup_{y \in \tilde{Y}} R(y)$. The set of vertices of $G$ is defined to be $U \cup V$.
In order to define the edges, consider any constraint $C=C(x, y) \in \tilde{\mathcal{C}}$. We define a set $E(C)$ of at most $d(n) \cdot A$ edges corresponding to $C$, as follows: we add an edge between vertex $v(x, a)$ and vertex $v\left(y, a^{\prime}\right)$ to $E(C)$ if assignments $a$ to $x$ and $a^{\prime}$ to $y$ satisfy the constraint $C$. Since instance $\mathcal{I}$ is a $d(n)$-to- $d(n)$ instance of Bipartite 2-CSP, every vertex of $R(x) \cup R(y)$ is incident to at most $d(n)$ edges of $E(C)$. We then let $E(G)=\bigcup_{C \in \tilde{\mathcal{C}}} E(C)$.
Let $k_{1}=|\tilde{X}|$ and $k_{2}=|\tilde{Y}|$. We can then view graph $G$, together with parameters $k_{1}$ and $k_{2}$ as an instance of the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, $\operatorname{DkS}\left(G, k_{1}, k_{2}\right)$. Notice that $|V(G)| \leq|\tilde{\mathcal{C}}| \cdot A^{2} \leq|\mathcal{C}| \cdot A^{2} \leq \operatorname{size}(\mathcal{I}) \leq n$. We apply Algorithm $\mathcal{A}$ to instance $\operatorname{DkS}\left(G, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph, and we let $S$ be the solution that the algorithm returns. Denote $G^{\prime}=G[S]$, so the value of the solution is $\left|E\left(G^{\prime}\right)\right|$. Assume first that $\left|E\left(G^{\prime}\right)\right|<\frac{|\tilde{\mathcal{C}}|}{4 \alpha(n)}$. We use the
following observation to show that, in this case, the set $\tilde{\mathcal{C}}$ of constraints is bad.
Observation 3.15. If the set $\tilde{\mathcal{C}}$ of constraints is not bad, then the value of the optimal solution to instance $\operatorname{DkS}\left(G, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph is at least $\frac{|\tilde{\mathcal{C}}|}{4}$.

Proof: Assume that the set $\tilde{\mathcal{C}}$ of constraints is not bad. Then there is an assignment $f$ to variables of $\tilde{X} \cup \tilde{Y}$ that satisfies more than $\frac{|\tilde{\mathcal{C}}|}{4}$ constraints of $\tilde{\mathcal{C}}$. For each variable $z \in \tilde{X} \cup \tilde{Y}$, we denote the corresponding assignment by $f(z)$. Let $S^{\prime}$ be the set of vertices of $G$ that contains, for every variable $x \in \tilde{X}$, vertex $v(x, f(x))$, and for every variable $y \in \tilde{Y}$, vertex $v(y, f(y))$. Then $S^{\prime}$ is a valid solution to instance $\operatorname{DkS}\left(G, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph. Moreover, for every constraint $C \in \tilde{\mathcal{C}}$ that is satisfied by the assignment $f$, an edge of $E(C)$ must be contained in $G\left[S^{\prime}\right]$. Therefore, the value of solution $S^{\prime}$ is at least $\frac{|\tilde{\mathcal{C}}|}{4}$.
From Observation 3.15 , if the set $\tilde{\mathcal{C}}$ of constraints is not bad, then algorithm $\mathcal{A}$ must have returned a solution whose value is at least $\frac{|\tilde{\mathcal{C}}|}{4 \alpha(n)}$. Therefore, if the value of the solution $S$ that the algorithm returns is less than $\frac{|\tilde{\mathcal{C}}|}{4 \cdot \alpha(n)}$, then we terminate the algorithm, and return $\tilde{\mathcal{C}}$ as a bad set of constraints. Recall that, from Claim $3.13,|\tilde{\mathcal{C}}|=|E(\tilde{H})| \geq \frac{|E(H)|}{8 \log ^{2} n}=\frac{|\mathcal{C}|}{8 \log ^{2} n}$. From now on we assume that the value of the soution $S$ is at least $\frac{|\tilde{\mathcal{C}}|}{4 \alpha(n)}$. We will use the set $S$ of vertices of $G$, in order to either correctly certify that graph $H$ is $\beta^{3}$-good, or to compute a subgraph $H^{\prime} \subseteq H$ with the required properties. It will be convenient for us to further regularize graph $G^{\prime}$, which we do in the next step.

### 3.4.3 Step 3: Further Regularization

It would be convenient for us to further regularize graph $G^{\prime}$, by computing a subgraph $G^{\prime \prime} \subseteq G^{\prime}$, with $\left|E\left(G^{\prime \prime}\right)\right|$ roughly comparable to $\left|E\left(G^{\prime}\right)\right|$, that has the following additional properties. First, for all variables $x \in \tilde{X}$ with $R(x) \cap V\left(G^{\prime \prime}\right) \neq \emptyset$, the cardinalities of the sets $R(x) \cap V\left(G^{\prime \prime}\right)$ of vertices are roughly equal to each other (to within factor 2). Similarly, for all variables $y \in \tilde{Y}$ with $R(y) \cap V\left(G^{\prime \prime}\right) \neq \emptyset$, the cardinalities of the sets $R(y) \cap V\left(G^{\prime \prime}\right)$ of vertices are roughly equal to each other. Lastly, for all constraints $C \in \tilde{\mathcal{C}}$ with $E(C) \cap E\left(G^{\prime \prime}\right) \neq \emptyset$, the cardinalities of the sets $E(C) \cap E\left(G^{\prime \prime}\right)$ of edges are roughly equal to each other. In this step, we compute a subgraph $G^{\prime \prime} \subseteq G^{\prime}$ with all these properties. The algorithm is summarized in the following claim.

Claim 3.16. There is an algorithm with running time $O(\underset{\sim}{\operatorname{poly}}(n))$, that computes a subgraph $G^{\prime \prime} \subseteq$ $G^{\prime}$, subsets $X^{*} \subseteq \tilde{X}, Y^{*} \subseteq \tilde{Y}$ of variables, a subset $\mathcal{C}^{*} \subseteq \tilde{\mathcal{C}}$ of constraints, and integers $q, q^{\prime}, r \geq 1$, such that, if we denote, for every variable $z \in X \cup Y, R^{\prime}(z)=R(z) \cap V\left(G^{\prime \prime}\right)$, and for every constraint $C \in \mathcal{C}, E^{\prime}(C)=E(C) \cap E\left(G^{\prime \prime}\right)$, then the following hold:

- for every variable $x \in X^{*}, 2^{q} \leq\left|R^{\prime}(x)\right|<2^{q+1}$, and for $x \notin X^{*}, R^{\prime}(x)=\emptyset$;
- for every variable $y \in Y^{*}, 2^{q^{\prime}} \leq\left|R^{\prime}(y)\right|<2^{q^{\prime}+1}$, and for $y \notin Y^{*}, R^{\prime}(y)=\emptyset$;
- for every constraint $C \in \mathcal{C}^{*}, 2^{r} \leq\left|E^{\prime}(C)\right|<2^{r+1}$, and for $C \notin \mathcal{C}^{*}, E^{\prime}(C)=\emptyset$; and
- $\left|\mathcal{C}^{*}\right| \geq \frac{|\tilde{\mathcal{C}}|}{2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n}$.

Proof: The proof follows a standard regularization process. Let $E^{0}=E\left(G^{\prime}\right)$, so that $\left|E^{0}\right| \geq \frac{|\tilde{\mathcal{C}}|}{4 \cdot \alpha(n)}$.
Our first step regularizes the variables of $\tilde{X}$. We group the variables of $\tilde{X}$ into sets $J_{0}, J_{1}, \ldots, J_{\lceil\log A\rceil}$. For all $0 \leq i \leq\lceil\log A\rceil$, we let $J_{i}=\left\{x \in \tilde{X}\left|2^{i} \leq\left|R(x) \cap V\left(G^{\prime}\right)\right|<2^{i+1}\right\}\right.$. Note that, if $R(x) \cap V\left(G^{\prime}\right)=\emptyset$, then variable $x$ does not belong to any of the groups that we have defined. We also partition the edges of $E^{0}$ into groups $E_{0}, E_{1}, \ldots, E_{\lceil\log A\rceil}$, where for all $0 \leq i \leq\lceil\log A\rceil$, group $E_{i}$ contains all edges $e \in E^{0}$ that are incident to the vertices of $\bigcup_{x \in J_{i}}\left(R(x) \cap V\left(G^{\prime}\right)\right)$. It is easy to verify that $\left(E_{0}, \ldots, E_{\lceil\log A\rceil}\right)$ is indeed a partition of the set $E^{0}$ of edges. Therefore, there is an index $0 \leq q \leq\lceil\log A\rceil$ with $\left|E_{q}\right| \geq \frac{\left|E^{0}\right|}{2 \log A} \geq \frac{\left|E^{0}\right|}{2 \log n}$. We then set $X^{*}=J_{q}$. For each such variable $x \in X^{*}$, we set $R^{\prime}(x)=R(x) \cap V\left(G^{\prime}\right)$, and for each variable $x \in \tilde{X} \backslash X^{*}$, we set $R^{\prime}(x)=\emptyset$. We also let $E^{1}=E_{q}$. From the above discussion, $\left|E^{1}\right| \geq \frac{\left|E^{0}\right|}{2 \log n}$, and, for every variable $x \in X^{*}$, $2^{q} \leq\left|R^{\prime}(x)\right|<2^{q+1}$. Note that all edges of $E^{1}$ are incident to vertices of $\bigcup_{x \in X^{*}} R^{\prime}(x)$.
Our second step is to regularize the variables of $\tilde{Y}$, exactly as before. We group the variables of $\tilde{Y}$ into sets $J_{0}^{\prime}, J_{1}^{\prime}, \ldots, J_{\lceil\log A\rceil}^{\prime}$. For all $0 \leq i^{\prime} \leq\lceil\log A\rceil$, we let $J_{i^{\prime}}^{\prime}=\left\{y \in \tilde{Y}\left|2^{i^{\prime}} \leq\right| R(x) \cap\right.$ $\left.V\left(G^{\prime}\right) \mid<2^{i^{\prime}+1}\right\}$. As before, if $R(y) \cap V\left(G^{\prime}\right)=\emptyset$, then variable $y$ does not belong to any of the sets that we have defined. We also partition the edges of $E^{1}$ into sets $E_{0}^{\prime}, E_{1}^{\prime}, \ldots, E_{\lceil\log A\rceil}^{\prime}$, where for all $0 \leq i^{\prime} \leq\lceil\log A\rceil$, set $E_{i^{\prime}}$ contains all edges $e \in E^{1}$ that are incident to the vertices of $\bigcup_{y \in J_{i^{\prime}}}\left(R(y) \cap V\left(G^{\prime}\right)\right)$. As before, $\left(E_{0}^{\prime}, \ldots, E_{\lceil\log A\rceil}^{\prime}\right)$ is a partition of the set $E^{1}$ of edges. Therefore, there is an index $0 \leq q^{\prime} \leq\lceil\log A\rceil$ with $\left|E_{q^{\prime}}^{\prime}\right| \geq \frac{\left|E^{1}\right|}{2 \log A} \geq \frac{\left|E^{1}\right|}{2 \log n} \geq \frac{\left|E^{0}\right|}{4 \log ^{2} n}$. We then set $Y^{*}=J_{q^{\prime}}^{\prime}$. For each variable $y \in Y^{*}$, we let $R^{\prime}(y)=R(y) \cap V\left(G^{\prime}\right)$, and for each variable $y \in \tilde{Y} \backslash Y^{*}$, we set $R^{\prime}(y)=\emptyset$. We also let $E^{2}=E_{q^{\prime}}^{\prime}$. From the above discussion, $\left|E^{2}\right| \geq \frac{\left|E^{0}\right|}{4 \log ^{2} n}$, and, for every variable $y \in Y^{*}, 2^{q^{\prime}} \leq\left|R^{\prime}(y)\right|<2^{y+1}$. Notice that every edge of $E^{2}$ is incident to a vertex of $\bigcup_{x \in X^{*}} R^{\prime}(x)$ and a vertex of $\bigcup_{y \in Y^{*}} R^{\prime}(y)$.
Our third and final step is to regularize the constraints. Recall that for each constraint $C=$ $C(x, y) \in \tilde{\mathcal{C}}$, we defined a set $E(C)$ of edges. Since $|R(x)|=A$ and $|R(y)|=A,|E(C)| \leq A^{2} \leq n$ must hold. We group all constraints $C \in \tilde{\mathcal{C}}$ into sets $\mathcal{C}_{0}, \mathcal{C}_{1}, \ldots, \mathcal{C}_{\lceil\log n\rceil}$, where for all $0 \leq j \leq\lceil\log n\rceil$, set $\mathcal{C}_{j}$ contains all constraints $C \in \tilde{\mathcal{C}}$ with $2^{j} \leq\left|E(C) \cap E^{2}\right|<2^{j+1}$. Note that, if $E(C) \cap E^{2}=\emptyset$, then constraint $C$ does not belong to any set. Next, we define a partition $E_{0}^{\prime \prime}, E_{1}^{\prime \prime}, \ldots, E_{\lceil\log n\rceil}^{\prime \prime}$ of the set $E^{2}$ of edges: for $0 \leq j \leq\lceil\log n\rceil$, set $E_{j}^{\prime \prime}$ contains all edges in $\bigcup_{C \in \mathcal{C}_{j}}\left(E(C) \cap E^{2}\right)$. It is easy to verify that $\left(E_{0}^{\prime \prime}, \ldots, E_{[\log n\rceil}^{\prime \prime}\right)$ is indeed a partition of $E^{2}$. Therefore, there must be an index $0<r \leq\lceil\log n\rceil$, with $\left|E_{r}^{\prime \prime}\right| \geq \frac{\left|E^{2}\right|}{2 \log n} \geq \frac{\left|E^{0}\right|}{8 \log ^{3} n}$. We let $\mathcal{C}^{*}=\mathcal{C}_{r}$ and $E^{3}=E_{r}^{\prime \prime}$. Since every constraint $C \in \mathcal{C}^{*}$ contributes at most $2^{r+1}$ edges to $E^{3}$, we get that:

$$
\left|\mathcal{C}^{*}\right| \geq \frac{\left|E^{3}\right|}{2^{r+1}} \geq \frac{\left|E^{0}\right|}{2^{r+4} \log ^{3} n} \geq \frac{|\tilde{\mathcal{C}}|}{2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n} .
$$

For every constraint $C \in \mathcal{C}^{*}$, we let $E^{\prime}(C)=E(C) \cap E^{2}$, and for every constraint $C \in \tilde{\mathcal{C}} \backslash \mathcal{C}^{*}$, we let $E^{\prime}(C)=\emptyset$. We are now ready to define graph $G^{\prime \prime}$. Its vertex set is $\left(\bigcup_{x \in X^{*}} R^{\prime}(X)\right) \cup\left(\bigcup_{y \in Y^{*}} R^{\prime}(y)\right)$, and its edge set is $\bigcup_{C \in \mathcal{C}^{*}} E^{\prime}(C)=E^{3}$. It is immediate to verify, from the above discussion, that graph $G^{\prime}$, sets $X^{*}, Y^{*}$ of variables, and set $\mathcal{C}^{*}$ of constraints have all required properties.

In the next observation, we establish some useful bounds on the cardinalities of the sets $X^{*}, Y^{*}$ of variables, and the set $\mathcal{C}^{*}$ of constraints.

Observation 3.17. All of the following bounds hold:

- $\left|X^{*}\right| \leq \frac{|\tilde{X}|}{2^{q}}$;
- $\left|Y^{*}\right| \leq \frac{|\tilde{Y}|}{2^{q}}$;
- $\left|\mathcal{C}^{*}\right| \leq 2 d_{1}\left|X^{*}\right| \leq \frac{2 d_{1}|\tilde{X}|}{2^{q}}$; and
- $\left|\mathcal{C}^{*}\right| \leq 2 d_{2}\left|Y^{*}\right| \leq \frac{2 d_{2}|\tilde{Y}|}{2^{q^{\prime}}}$.

Proof: From the definition of the instance ( $G, k_{1}, k_{2}$ ) of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph, graph $G^{\prime}$ contains at most $k_{1}=|\tilde{X}|$ vertices of $\bigcup_{x \in \tilde{X}} R(x)$, and at most $k_{2}=|\tilde{Y}|$ vertices of $\bigcup_{y \in \tilde{Y}} R(y)$. Since, for every variable $x \in X^{*}, R^{\prime}(x) \subseteq V\left(G^{\prime}\right)$ and $\left|R^{\prime}(x)\right| \geq 2^{q}$, we get that $\left|X^{*}\right| \leq \frac{|\tilde{X}|}{2^{q}}$. Similarly, $\left|Y^{*}\right| \leq \frac{|\tilde{Y}|}{2 q^{\prime}}$.
Recall that, since $\tilde{H}$ is a nice subgraph of $H$, for every vertex $x \in \tilde{H}, \operatorname{deg}_{\tilde{H}}(x) \leq \operatorname{deg}_{H}(x) \leq 2 d_{1}$, and so $x$ may participate in at most $2 d_{1}$ constraints of $\tilde{\mathcal{C}}$. Since $\mathcal{C}^{*} \subseteq \tilde{\mathcal{C}}$ and $X^{*} \subseteq \tilde{X}$, every variable $x \in X^{*}$ may participate in at most $2 d_{1}$ constraints of $\mathcal{C}^{*}$. Therefore, $\left|\mathcal{C}^{*}\right| \leq 2 d_{1}\left|X^{*}\right|$.
Similarly, since $\tilde{H}$ is a nice subgraph of $H$, for every vertex $y \in \tilde{Y}$, $\operatorname{deg}_{\tilde{H}}(y) \leq 2 d_{2}$. Using the same arguments as before, $\left|\mathcal{C}^{*}\right| \leq 2 d_{2}\left|Y^{*}\right|$.
Recall that, from Claim $3.16,\left|\mathcal{C}^{*}\right| \geq \frac{|\tilde{\mathcal{C}}|}{2^{r+5} \cdot \alpha(n) \cdot \log ^{3} n}$. Since graph $\tilde{H}$ is a nice subgraph of $H$, we get that, for every vertex $y \in \tilde{Y}, \operatorname{deg}_{\tilde{H}}(y) \geq d_{2}$. Therefore, $|\tilde{\mathcal{C}}| \geq|\tilde{Y}| \cdot d_{2}$, and so:

$$
\begin{equation*}
\left|\mathcal{C}^{*}\right| \geq \frac{d_{2} \cdot|\tilde{Y}|}{2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n} \tag{1}
\end{equation*}
$$

Similarly, from the definition of a nice subgraph, $|\tilde{\mathcal{C}}| \geq \frac{d_{1}}{4 \log n} \cdot|\tilde{X}|$, and so:

$$
\begin{equation*}
\left|\mathcal{C}^{*}\right| \geq \frac{d_{1} \cdot|\tilde{X}|}{2^{r+8} \cdot \alpha(n) \cdot \log ^{4} n} . \tag{2}
\end{equation*}
$$

Lastly, we show that both $2^{q}, 2^{q^{\prime}}$ are close to $2^{r}$, in the following corollary of Observation 3.17.
Corollary 3.18. The following inequalities hold:

- $\frac{2^{r}}{2 d(n)} \leq 2^{q} \leq 2^{r+8} \cdot \alpha(n) \cdot \log ^{4} n$; and
- $\frac{2^{r}}{2 d(n)} \leq 2^{q^{\prime}} \leq 2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n$.

Proof: Consider some constraint $C=C(x, y) \in \mathcal{C}^{*}$, and recall that $\left|E^{\prime}(C)\right| \geq 2^{r}$. From the definition of the $d$-to- $d$ instances, every vertex $v(x, a) \in R(x)$ may be incident to at most $d(n)$ edges of $E(C)$. Since all edges of $E^{\prime}(C)$ are incident to vertices of $R^{\prime}(x)$, and $\left|R^{\prime}(x)\right| \leq 2^{q+1}$, we get that $\left|E^{\prime}(C)\right| \leq d(n) \cdot\left|R^{\prime}(x)\right| \leq d(n) \cdot 2^{q+1}$. We conclude that $2^{r} \leq d(n) \cdot 2^{q+1}$.
Similarly, every vertex $v\left(y, a^{\prime}\right) \in R(y)$ may be incident to at most $d(n)$ edges of $E(C)$. Since all edges of $E^{\prime}(C)$ are incident to vertices of $R^{\prime}(y)$, and $\left|R^{\prime}(y)\right| \leq 2^{q^{\prime}+1}$, we get that $\left|E^{\prime}(C)\right| \leq d(n) \cdot 2^{q^{\prime}+1}$. This proves the inequalities $\frac{2^{r}}{2 d(n)} \leq 2^{q}$ and $\frac{2^{r}}{2 d(n)} \leq 2^{q^{\prime}}$.
Next, we prove that $2^{q^{\prime}} \leq 2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n$. Recall that, from Inequality $1,\left|\mathcal{C}^{*}\right| \geq \frac{d_{2} \cdot|\tilde{Y}|}{2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n}$. On the other hand, from the definition of a nice subgraph, every variable $y \in Y^{*}$ may participate in at most $2 d_{2}$ constraints of $\mathcal{C}^{*}$, and, from Observation $3.17,\left|Y^{*}\right| \leq \frac{|\tilde{Y}|}{2 q^{\prime}}$. Therefore:

$$
\begin{equation*}
\left|\mathcal{C}^{*}\right| \leq 2 d_{2} \cdot\left|Y^{*}\right| \leq \frac{2 d_{2}|\tilde{Y}|}{2^{q^{\prime}}} \tag{3}
\end{equation*}
$$

Combining the two inequalities, we get that: $2^{q^{\prime}} \leq 2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n$.
Lastly, we prove that $2^{q} \leq 2^{r+8} \cdot \alpha(n) \cdot \log ^{4} n$. Recall that, from Inequality $\left|2,\left|\mathcal{C}^{*}\right| \geq \frac{d_{1} \cdot|\tilde{X}|}{2^{r+7} \cdot \alpha(n) \cdot \log ^{4} n}\right.$ holds. As before, from the definition of a nice subgraph, every variable $x \in X^{*}$ may participate in at most $2 d_{1}$ constraints of $\mathcal{C}^{*}$, and, from Observation 3.17. $\left|X^{*}\right| \leq \frac{|\tilde{X}|}{2^{q}}$. Therefore:

$$
\begin{equation*}
\left|\mathcal{C}^{*}\right| \leq 2 d_{1} \cdot\left|X^{*}\right| \leq \frac{2 d_{1} \cdot|\tilde{X}|}{2^{q}} . \tag{4}
\end{equation*}
$$

Combining the two inequalities together, we get that: $2^{q} \leq 2^{r+8} \cdot \alpha(n) \cdot \log ^{4} n$.

### 3.4.4 Step 4: Certifying that $H$ is a Good Graph or Computing a Subgraph of $H$

We consider two cases, depending on whether $2^{r} \leq \beta$ holds. We start by showing that, if $2^{r} \leq \beta$, then graph $H$ is $\beta^{3}$-good.

Observation 3.19. If $2^{r} \leq \beta$, then graph $H$ is $\beta^{3}$-good.
Proof: We show that there exists an assignment to variables of $X \cup Y$ that satisfies at least $|\mathcal{C}| / \beta^{3}$ constraints of $\mathcal{C}$. In order to do it, we show a randomized algorithm that computes assignments to variables of $X \cup Y$, such that the expected number of satisfied constraints is at least $|\mathcal{C}| / \beta^{3}$.

The assignments are computed as follows. Consider a variable $x \in X$. If $x \notin X^{*}$, then we assign to $x$ an arbitrary value from $[A]$. Assume now that $x \in X^{*}$. Recall that we have defined a set $R^{\prime}(x) \subseteq R(x)$ of vertices, whose cardinality is at most $2^{q+1}$. Set $R^{\prime}(x)$ of vertices naturally defines a collection $\hat{A}(x)=\left\{a \in[A] \mid v(x, a) \in R^{\prime}(x)\right\}$ of assignments to variable $x$, with $|\hat{A}(x)| \leq 2^{q+1}$. We choose an assignment $a \in \hat{A}(x)$ uniformly at random, and assign value $a$ to $x$.

Assignments to variables of $Y$ are defined similarly. Consider a variable $y \in Y$. If $y \notin Y^{*}$, then we assign to $y$ an arbitrary value from $[A]$. Assume now that $y \in Y^{*}$. Recall that we have defined a set $R^{\prime}(y) \subseteq R(y)$ of vertices, whose cardinality is at most $2^{q^{\prime}+1}$. Set $R^{\prime}(y)$ of vertices naturally defines
a collection $\hat{A}(y)=\left\{a \in[A] \mid v(y, a) \in R^{\prime}(y)\right\}$ of assignments to variable $y$, with $|\hat{A}(y)| \leq 2^{q^{\prime}+1}$. We choose an assignment $a^{\prime} \in \hat{A}(y)$ uniformly at random, and assign value $a^{\prime}$ to $y$.
Recall that we have computed, in Claim $\sqrt{3.16}$, a collection $\mathcal{C}^{*} \subseteq \mathcal{C}$ of constraints, with $\left|\mathcal{C}^{*}\right| \geq$ $\frac{|\tilde{\mathcal{C}}|}{2^{r+6 \cdot \alpha(n) \cdot \log ^{3} n}}$. Consider now any constraint $C=C(x, y) \in \mathcal{C}^{*}$, and recall that $x \in X^{*}, y \in Y^{*}$ must hold. Recall that graph $G^{\prime \prime}$ contains a collection $E^{\prime}(C) \subseteq E(C)$ of edges, with $\left|E^{\prime}(C)\right| \geq 2^{r}$. Consider now any such edge $e=(v(x, a), v(y, A))$. We say that edge $e$ wins if $x$ is assigned value $a$, and $y$ is assigned value $a^{\prime}$. The probability that edge $e$ wins is at least $\frac{1}{2^{q+1} \cdot 2^{q^{\prime}+1}}$. Notice that at most one edge of $E^{\prime}(C)$ may win, and so the probability that any edge of $E^{\prime}(C)$ wins is at least $\frac{\left|E^{\prime}(C)\right|}{2^{q+1} \cdot 2^{q^{\prime}+1}} \geq \frac{2^{r}}{2^{q+1} \cdot 2^{q^{\prime}+1}}$. If any edge of $E^{\prime}(C)$ wins, the constraint $C$ is satisfied by the assignment that the algorithm chooses. Therefore, the probability that a constraint $C \in \mathcal{C}^{*}$ is satisfied is at least $\frac{2^{r}}{2^{q+1} \cdot 2^{q^{\prime}+1}}$.
Overall, the expected number of constraints that are satisfied by the assignment is at least:

$$
\frac{\left|\mathcal{C}^{*}\right| \cdot 2^{r}}{2^{q+1} \cdot 2^{q^{\prime}+1}} \geq \frac{|\tilde{\mathcal{C}}|}{2^{8} \cdot \alpha(n) \cdot 2^{q} \cdot 2^{q^{\prime}} \cdot \log ^{3} n}
$$

Recall that, from Claim $3.13,|\tilde{\mathcal{C}}|=|E(\tilde{H})| \geq \frac{|E(H)|}{8 \log ^{2} n}=\frac{|\mathcal{C}|}{8 \log ^{2} n}$, and, from Corollary $3.18,2^{q} \cdot 2^{q^{\prime}} \leq$ $2^{2 r+14} \cdot(\alpha(n))^{2} \cdot \log ^{7} n \leq 2^{14} \cdot \beta^{2} \cdot(\alpha(n))^{2} \cdot \log ^{7} n$, since we have assumed that $2^{r} \leq \beta$. Therefore, the expected number of constraints of $\mathcal{C}$ that are satisfied by the assignment is at least:

$$
\frac{|\tilde{\mathcal{C}}|}{2^{8} \cdot \alpha(n) \cdot 2^{q} \cdot 2^{q^{\prime}} \cdot \log ^{3} n} \geq \frac{|\mathcal{C}|}{2^{11} \cdot \alpha(n) \cdot 2^{q} \cdot 2^{q^{\prime}} \cdot \log ^{5} n} \geq \frac{|\mathcal{C}|}{2^{25} \cdot \beta^{2} \cdot(\alpha(n))^{3} \cdot \log ^{12} n} \geq \frac{|\mathcal{C}|}{\beta^{3}},
$$

since $\beta \geq 2^{27}(\alpha(n))^{3} \log ^{12} n$.
We conclude that there is an assignment to the variables of $X \cup Y$ that satisfies at least $|\mathcal{C}| / \beta^{3}$ constraints of $\mathcal{C}$, and so graph $H$ is $\beta^{3}$-good.
If $2^{r} \leq \beta$, then we terminate the algorithm and report that graph $H$ is $\beta$-good.
From now on we assume that $2^{r}>\beta$. In this case, we return a subgraph a subgraph $H^{\prime}=$ $\left(X^{\prime}, Y^{\prime}, E^{\prime}\right)$ of $H(\mathcal{I})$, that is defined as follows: $X^{\prime}=X^{*}, Y^{\prime}=Y^{*}$, and $E^{\prime}=E^{*}$. We now verify that this graph has all required properties.
Recall that, from Observation $3.17,\left|X^{*}\right| \leq \frac{|\tilde{X}|}{2^{q}} \leq \frac{|X|}{2^{q}}$, from Corollary 3.18, $2^{q} \geq \frac{2^{r}}{2 d(n)}$, and, from our assumption, $2^{r}>\beta$. Therefore:

$$
\left|X^{*}\right| \leq \frac{|X|}{2^{q}} \leq \frac{|X| \cdot 2 d(n)}{2^{r}} \leq \frac{2 d(n)}{\beta} \cdot|X| .
$$

Similarly, from Observation 3.17. $\left|Y^{*}\right| \leq \frac{|\tilde{Y}|}{2^{q^{\prime}}} \leq \frac{|Y|}{2^{q^{\prime}}}$, from Corollary 3.18. $2^{q^{\prime}} \geq \frac{2^{r}}{2 d(n)}$, and, from our assumption, $2^{r}>\beta$. Therefore:

$$
\left|Y^{*}\right| \leq \frac{|Y|}{2^{q^{\prime}}} \leq \frac{|Y| \cdot 2 d(n)}{2^{r}} \leq \frac{2 d(n)}{\beta} \cdot|Y| .
$$

It now only remains to show that $\left|E^{*}\right| \geq \frac{\operatorname{vol}_{H}\left(X^{*} \cup Y^{*}\right)}{256 d(n) \cdot \alpha(n) \cdot \log ^{4} n}$.
Recall that, from Inequality $\left|2,\left|\mathcal{C}^{*}\right| \geq \frac{d_{1} \cdot|\tilde{X}|}{2^{r+8} \cdot \alpha(n) \cdot \log ^{4} n}\right.$ holds. Since, from Corollary $3.18,2^{r} \leq 2^{q} \cdot 2 d(n)$, we get that $\left|E^{*}\right|=\left|\mathcal{C}^{*}\right| \geq \frac{d_{1} \cdot|\tilde{X}|}{2^{q+9} \cdot d(n) \cdot \alpha(n) \cdot \log ^{4} n}$. At the same time, from the definition of a nice graph, for every vertex $x \in X^{*}, \operatorname{deg}_{H}(x) \leq 2 d_{1}$, so $\operatorname{vol}_{H}\left(X^{*}\right) \leq 2 d_{1} \cdot\left|X^{*}\right| \leq \frac{2 d_{1} \cdot|\tilde{X}|}{2^{q}}$, since $\left|X^{*}\right| \leq \frac{|\tilde{X}|}{2^{q}}$ from Observation 3.17. Therefore, $\left|E^{\prime}\right| \geq \frac{\operatorname{vol}_{H}\left(X^{*}\right)}{1024 d(n) \cdot \alpha(n) \cdot \log ^{4} n}$.
Similarly, from Inequality $11,\left|\mathcal{C}^{*}\right| \geq \frac{d_{2} \cdot|\tilde{Y}|}{2^{r+6} \cdot \alpha(n) \cdot \log ^{3} n}$. Since, from Corollary $3.18,2^{r} \leq 2^{q^{\prime}} \cdot 2 d(n)$, we get that $\left|E^{\prime}\right|=\left|\mathcal{C}^{*}\right| \geq \frac{d_{2} \cdot|\tilde{Y}|}{2^{q^{\prime}+7} \cdot d(n) \cdot \alpha(n) \cdot \log ^{3} n}$. At the same time, from the definition of a nice graph, for every vertex $y \in Y^{*}, \operatorname{deg}_{H}(y) \leq 2 d_{2}$, so $\operatorname{vol}_{H}\left(Y^{*}\right) \leq 2 d_{2} \cdot\left|Y^{*}\right| \leq \frac{2 d_{2} \cdot \mid \tilde{Y}}{2^{q^{\prime}}}$, since $\left|Y^{*}\right| \leq \frac{|\tilde{Y}|}{2^{q^{\prime}}}$ from Observation 3.17. Therefore, $\left|E^{*}\right| \geq \frac{\operatorname{vol}_{H}\left(Y^{*}\right)}{1024 d(n) \cdot \alpha(n) \cdot \log ^{3} n}$.
Altogether, we get that $\left|E^{\prime}\right| \geq \frac{\operatorname{vol}_{H}\left(X^{*} \cup Y^{*}\right)}{2048 d(n) \cdot \alpha(n) \cdot \log ^{4} n}$.
Note that every step of the algorithm, except Step 2, has running time $O(\operatorname{poly}(n))$, while the running time of Step 2 is $O(T(n)+\operatorname{poly}(n))$. Therefore, the total running time of the algorithm is at most $O(T(n) \cdot \operatorname{poly}(n))$.

## 4 Reductions from Dense $k$-Coloring and ( $r, h$ )-Graph Partitioning to Densest $k$-Subgraph

In this section we provide reductions from the Dense $k$-Coloring and ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning problems to Densest $k$-Subgraph, by proving the following theorem.

Theorem 4.1. Let $\alpha: \mathbb{Z}^{+} \rightarrow \mathbb{Z}^{+}$be an increasing function, such that $\alpha(n) \leq o(n)$. Assume that there is an efficient $\alpha(n)$-approximation algorithm for the Densest $k$-Subgraph problem, where $n$ is the number of vertices in the input graph. Then both of the following hold:

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

We provide the proof of the first assertion of the theorem, by showing a reduction from Dense $k$-Coloring to Densest $k$-Subgraph. The proof of the second assertion is similar and is deferred to Section B of Appendix. We start by considering an LP-relaxation of the Dense $k$-Coloring problem, whose number of variables is at least $\binom{N}{k}$. Due to this high number of variables, we cannot solve it directly. We first show an algorithm, that, given an approximate fractional solution to this LP-relaxation, whose support size is polynomial in $N$, computes an approximate integral solution
to the Dense $k$-Coloring problem instance. We then show an efficient algorithm that computes an approximate solution to the LP-relaxation, whose support is relatively small. In order to do so, we design an approximate separation oracle to the dual LP of the LP-relaxation.

### 4.1 An LP-Relaxation and Its Rounding

Let $\operatorname{DkC}(G, k)$ be an input instance of the Dense $k$-Coloring problem, and denote $|V(G)|=N$. We let $\mathcal{R}$ the collection of all subsets of $V(G)$ containing at most $k$ vertices, that is: $\mathcal{R}=\{S \subseteq$ $V(G)||S| \leq k\}$. For every set $S \in \mathcal{R}$ of vertices, we denote by $m(S)=\left|E_{G}(S)\right|$. We consider the following LP-relaxation of the Dense $k$-Coloring problem, that has a variable $x_{S}$ for every set $S \in \mathcal{R}$ of vertices.

$$
\begin{array}{rcl}
(\text { LP-P) } & \\
\max & \sum_{S \in \mathcal{R}} m(S) \cdot x_{S} & \\
\text { s.t. } & & \\
& \sum_{\substack{S \in \mathcal{R} \\
v \in S}} x_{S} \leq 1 & \forall v \in V(G) \\
& \sum_{S \in \mathcal{R}} x_{S} \leq N / k & \\
& x_{S} \geq 0 \quad \forall S \in \mathcal{R}
\end{array}
$$

It is easy to verify that (LP-P) is an LP-relaxation of the Dense $k$-Coloring problem. Indeed, consider a solution $\left(S_{1}, \ldots, S_{r}\right)$ to the input instance $\operatorname{DkC}(G, k)$, where $r=N / k$. For all $1 \leq i \leq r$, we set $x_{S_{i}}=1$, and for every other set $S \in \mathcal{R}$, we set $x_{S}=0$. This provides a feasible solution to (LP-P), whose value is precisely $\sum_{i=1}^{r}\left|E_{G}\left(S_{i}\right)\right|$. We denote the value of the optimal solution to (LP-P) by $\mathrm{OPT}_{\mathrm{LP}-\mathrm{P}}$. From the above discussion, $\mathrm{OPT}_{\mathrm{LP}-\mathrm{P}} \geq \mathrm{OPT}_{\mathrm{DkC}}(G, k)$.
Note that the number of variables in (LP-P) is at least $\binom{N}{k}$, and so we cannot solve it directly. We will show below an efficient algorithm that provides an approximate solution to (LP-P), whose support is reasonably small. Before we do so, we provide an LP-rounding algorithm, by proving the following claim.

Claim 4.2. There is an efficient randomized algorithm, whose input consists of an instance $\operatorname{DkC}(G, k)$ of the Dense $k$-Coloring problem with $N=|V(G)|$, such that $N$ is greater than a large enough constant, and a solution $\left\{x_{S} \mid S \in \mathcal{R}\right\}$ to (LP-P), in which the number of variables $x_{S}$ with $x_{S}>0$ is bounded by $\operatorname{poly}(N)$, and $\sum_{S \in \mathcal{R}} m(S) \cdot x_{S} \geq O P T_{\mathrm{LP}-\mathrm{P}} / \beta$, for some parameter $1 \leq \beta \leq N^{3}$; the solution is given by only specifying values of variables $x_{S}$ that are non-zero. The algorithm with high probability returns an integral solution $\left(S_{1}, \ldots, S_{N / k}\right)$ to instance $\operatorname{DkC}(G, k)$, such that $\sum_{i=1}^{N / k}\left|E_{G}\left(S_{i}\right)\right| \geq \frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{2000 \beta \log ^{3} N}$.

Proof: We assume that we are given a solution $\left\{x_{S} \mid S \in \mathcal{R}\right\}$ to (LP-P), in which the number of variables $x_{S}$ with $x_{S}>0$ is bounded by poly $(N)$. Denote $C=\sum_{S \in \mathcal{R}} m(S) \cdot x_{S}$, so that $C \geq \frac{\mathrm{OPT}_{\mathrm{LP}-\mathrm{P}}}{\beta} \geq \frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{\beta}$ holds. We denote by $\mathcal{R}^{\prime} \subseteq \mathcal{R}$ the collection of all sets $S \in \mathcal{R}$ with $x_{S}>0$. Assume first that there is any set $S \in \mathcal{R}^{\prime}$ with $m(S) \geq \frac{C}{300 \log ^{3} N}$. Then we can obtain the
desired solution $\left(S_{1}, \ldots, S_{N / k}\right)$ as follows. We start with $S_{1}=S$ and $S_{2}=\cdots=S_{N / k}=\emptyset$. We then iteratively add vertices of $V(G) \backslash S$ to sets $S_{1}, \ldots, S_{N / k}$ arbitrarily, to ensure that the cardinality of each such set is exactly $k$. It is immediate to verify that $\sum_{i=1}^{N / k}\left|E_{G}\left(S_{i}\right)\right| \geq \frac{C}{300 \log ^{3} N} \geq \frac{\mathrm{OPT}_{\mathrm{DKC}}(G, k)}{300 \beta \log ^{3} N}$. Therefore, we assume from now on that, for every set $S \in \mathcal{R}^{\prime}, m(S)<\frac{C}{300 \log ^{3} N}$.
We construct another collection $\mathcal{R}^{\prime \prime} \subseteq \mathcal{R}^{\prime}$ of subsets of vertices of $G$ as follows. For every set $S \in \mathcal{R}^{\prime}$, we add $S$ to $\mathcal{R}^{\prime \prime}$ independently, with probability $x_{S}$. Clearly, $\mathbf{E}\left[\sum_{S \in \mathcal{R}^{\prime \prime}} m(S)\right]=\sum_{S \in \mathcal{R}} m(S) \cdot x_{S}=$ $C$.

We say that a bad event $\mathcal{E}_{1}$ happens if some vertex $v \in V(G)$ lies in more than $5 \log N$ sets in $\mathcal{R}^{\prime \prime}$. We say that a bad event $\mathcal{E}_{2}$ happens if $\left|\mathcal{R}^{\prime \prime}\right|>(5 N \log N) / k$. We say that a bad event $\mathcal{E}_{3}$ happens if $\sum_{S \in \mathcal{R}^{\prime \prime}} m(S)<\frac{C}{8}$. Lastly, we say that a bad event $\mathcal{E}$ happens if either of the events $\mathcal{E}_{1}, \mathcal{E}_{2}$, or $\mathcal{E}_{3}$ happen. We start with the following simple observation.
Observation 4.3. $\operatorname{Pr}[\mathcal{E}] \leq \frac{2}{N^{3}}$.
Proof: For every set $S \in \mathcal{R}^{\prime}$, let $Y_{S}$ be the random variable whose value is 1 if $S \in \mathcal{R}^{\prime \prime}$ and 0 otherwise.

Consider any vertex $v \in V(G)$. Denote by $Z_{v}$ the number of vertex sets in $\mathcal{R}^{\prime \prime}$ containing $v$. Clearly, $Z_{v}=\sum_{\substack{s \in \mathcal{R}^{\prime},: \\ v \in S}} Y_{S}$. Therefore:

$$
\mathbf{E}\left[Z_{v}\right]=\mathbf{E}\left[\sum_{\substack{S \in \mathcal{R}^{\prime}: \\ v \in S}} Y_{S}\right]=\sum_{\substack{S \in \mathcal{R}^{\prime}: \\ v \in S}} x_{S} \leq 1,
$$

where the last inequality follows from the constraints of (LP-P).
By applying the Chernoff Bound from Lemma 2.2, we get that $\operatorname{Pr}\left[Z_{v}>5 \log N\right] \leq 1 / N^{5}$. Using the union bound over all vertices $v \in V(G)$, we get that $\operatorname{Pr}\left[\mathcal{E}_{1}\right] \leq 1 / N^{4}$.
Notice that $\mathbf{E}\left[\left|\mathcal{R}^{\prime \prime}\right|\right]=\mathbf{E}\left[\sum_{S \in \mathcal{R}} Y_{S}\right]=\sum_{S \in \mathcal{R}} x_{S} \leq N / k$ from the constraints of (LP-P). Bad Event $\mathcal{E}_{2}$ happens if $\left|\mathcal{R}^{\prime \prime}\right|=\sum_{S \in \mathcal{R}} Y_{S}>(5 N \log N) / k$. By applying the Chernoff Bound from Lemma 2.2 to the variables of $\left\{Y_{S} \mid S \in \mathcal{R}^{\prime}\right\}$, we get that $\operatorname{Pr}\left[\mathcal{E}_{2}\right]=\operatorname{Pr}\left[\sum_{S \in \mathcal{R}^{\prime}} Y_{S}>(5 N \log N) / k\right] \leq 1 / N^{5}$.

Lastly, we bound the probability of Event $\mathcal{E}_{3}$. Recall that we have assumed that, for every set $S \in \mathcal{R}^{\prime}, m(S)<\frac{C}{300 \log ^{3} N}$ holds. We partition the collection $\mathcal{R}^{\prime}$ of vertex subsets into $\rho=\lceil 2 \log N\rceil$ collections $\mathcal{R}_{0}, \ldots, \mathcal{R}_{\rho-1}$, as follows. For all $0 \leq i<\rho$, we let $\mathcal{R}_{i}=\left\{S \in \mathcal{R}^{\prime} \mid 2^{i} \leq m(S)<2^{i+1}\right\}$. For all $0 \leq i<\rho$, we denote $C_{i}=\sum_{S \in \mathcal{R}_{i}} m(S) \cdot x_{S}$. Clearly, $\sum_{i=0}^{\rho-1} C_{i}=C$. We say that an index $0 \leq i<\rho$ is bad, if $C_{i}<\frac{C}{8 \log N}$, and otherwise we say that $i$ is a good index. Let $I^{g}, I^{b} \subseteq\{0, \ldots, \rho-1\}$ be the collections of good and bad indices, respectively. Since $\rho \leq 4 \log N$, we get that:

$$
\sum_{i \in I^{b}} C_{i} \leq(4 \log N) \cdot \frac{C}{8 \log N} \leq \frac{C}{2}
$$

Therefore, $\sum_{C} i \in I^{g} C_{i} \geq \frac{C}{2}$ holds. Consider now a good index $i \in I^{g}$. Recall that, for every set $S \in \mathcal{R}^{\prime}$, $m(S) \leq \frac{C}{300 \log ^{3} N}$ holds, and so $2^{i} \leq \frac{C}{300 \log ^{3} N}$ must hold. Moreover, since $\sum_{S \in \mathcal{R}_{i}} m(S) \cdot x_{S}=C_{i} \geq$ $\frac{C}{8 \log N}$, we get that:

$$
\frac{C}{8 \log N} \leq \sum_{S \in \mathcal{R}_{i}} m(S) x_{S} \leq 2^{i+1} \cdot \sum_{S \in \mathcal{R}_{i}} x_{S} \leq \frac{C}{150 \log ^{3} N} \cdot \sum_{S \in \mathcal{R}_{i}} x_{S}
$$

We conclude that $\sum_{S \in \mathcal{R}_{i}} x_{S} \geq 16 \log ^{2} N$ holds for every good index $i \in I^{g}$.
For a good index $i \in I^{g}$, we denote by $\mathcal{R}_{i}^{\prime}=\mathcal{R}_{i} \cap \mathcal{R}^{\prime \prime}$, and we let $\tilde{\mathcal{E}}_{i}$ be the bad event that $\left|\mathcal{R}_{i}^{\prime}\right|<\frac{\sum_{S \in \mathcal{R}_{i}} x_{S}}{2}$. From the Chernoff Bound (Lemma 2.2, and the fact that $\sum_{S \in \mathcal{R}_{i}} x_{S} \geq 16 \log ^{2} N$, we get that $\operatorname{Pr}\left[\tilde{\mathcal{E}}_{i}\right] \leq e^{-2 \log ^{2} N}<N^{-4}$, if $N$ is sufficiently large. By applying the Union Bound to all indices $i \in I^{g}$, we get that the probability that any of the events in $\left\{\tilde{\mathcal{E}}_{i} \mid i \in I^{g}\right\}$ happens is bounded by $1 / N^{3}$. Note that, if neither of the events in $\left\{\tilde{\mathcal{E}}_{i} \mid i \in I^{g}\right\}$ happen, then:

$$
\begin{aligned}
\sum_{S \in \mathcal{R}^{\prime \prime}} m(S) & \geq \sum_{i \in I^{g}} 2^{i} \cdot\left|\mathcal{R}_{i}^{\prime}\right| \\
& \geq \sum_{i \in I^{g}} 2^{i} \cdot \frac{\sum_{S \in \mathcal{R}_{i}} x_{S}}{2} \\
& \geq \sum_{i \in I^{g}} \sum_{S \in \mathcal{R}_{i}} \frac{m(S) \cdot x_{S}}{4} \\
& \geq \sum_{i \in I^{g}} \frac{C_{i}}{4} \\
& \geq \frac{C}{8}
\end{aligned}
$$

Therefore, if neither of the events in $\left\{\tilde{\mathcal{E}}_{i} \mid i \in I^{g}\right\}$ happen, then Event $\mathcal{E}_{3}$ also does not happen. We conclude that $\operatorname{Pr}\left[\mathcal{E}_{3}\right] \leq 1 / N^{3}$.
Finally, from the Union Bound, we get that:

$$
\operatorname{Pr}[\mathcal{E}] \leq \operatorname{Pr}\left[\mathcal{E}_{1}\right]+\operatorname{Pr}\left[\mathcal{E}_{2}\right]+\operatorname{Pr}\left[\mathcal{E}_{3}\right] \leq \frac{1}{N^{4}}+\frac{1}{N^{5}}+\frac{1}{N^{3}} \leq \frac{2}{N^{3}} .
$$

Observe that we can efficiently check whether Event $\mathcal{E}$ happened. If Event $\mathcal{E}$ happens, then we terminate the algorithm with a FAIL. We assume from now on that Event $\mathcal{E}$ did not happen. In this case, we are guaranteed that $\sum_{S \in \mathcal{R}^{\prime \prime}} m(S) \geq \frac{C}{8} \geq \frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{8 \beta}$. We denote $\mathcal{R}^{\prime \prime}=\left\{S_{1}, S_{2}, \ldots, S_{z}\right\}$, where the sets are indexed according to their value $m(S)$, so that $m\left(S_{1}\right) \geq m\left(S_{2}\right) \geq \cdots \geq m\left(S_{z}\right)$. We then let $\mathcal{S}=\left\{S_{1}, \ldots, S_{N / k}\right\}$ (if $z<N / k$, then we set $S_{z+1}=\cdots=S_{N / k}=\emptyset$ ). For all $1 \leq i \leq N / k$, we denote $E_{i}=E_{G}\left(S_{i}\right)$, so $\left|E_{i}\right|=m\left(S_{i}\right)$, and we denote $E^{\prime}=\bigcup_{i=1}^{N / k} E_{i}$. Recall that, since Event $\mathcal{E}$ did not happen, $\left|\mathcal{R}^{\prime \prime}\right| \leq(5 N \log N) / k$ holds. Therefore:

$$
\left|E^{\prime}\right| \geq \frac{\sum_{S \in \mathcal{R}^{\prime \prime}} m(S)}{5 \log N} \geq \frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{40 \beta \log N}
$$

Note that the vertex sets in the family $\mathcal{S}$ may not be mutually disjoint. However, since Event $\mathcal{E}$ did not happen, every vertex of $V(G)$ may lie in at most $5 \log N$ such sets. We now construct a new collection $\mathcal{S}^{\prime}=\left\{S_{1}^{\prime}, \ldots, S_{N / k}^{\prime}\right\}$ of sets of vertices, as follows. Consider any vertex $v \in V(G)$, and let $S_{i_{1}}, S_{i_{2}}, \ldots, S_{i_{a}} \in \mathcal{S}$ be the sets of $\mathcal{S}$ containing $v$. Vertex $v$ chooses an index $i^{*} \in\left\{i_{1}, \ldots, i_{a}\right\}$ at random, and is then added to $S_{i^{*}}^{\prime}$.
Note that for all $1 \leq j \leq N / k$, for every vertex $v \in S_{j}$, the probability that $v \in S_{j}^{\prime}$ is at least $1 /(5 \log N)$. We say that an edge $e=(u, v) \in E_{j}$ survives if both $u, v \in S_{j}^{\prime}$. We denote by $E^{\prime \prime} \subseteq E^{\prime}$ the set of all edges that survive. Since $\operatorname{Pr}\left[u \in S_{j}^{\prime}\right] \geq 1 /(5 \log N), \operatorname{Pr}\left[v \in S_{j}^{\prime}\right] \geq 1 /(5 \log N)$, and the two events are independent, we get that the probability that edge $e$ survives is at least $1 /\left(25 \log ^{2} N\right)$. Overall, we get that:

$$
\mathbf{E}\left[\left|E^{\prime \prime}\right|\right] \geq \frac{\left|E^{\prime}\right|}{25 \log ^{2} N} \geq \frac{\operatorname{OPT}_{\mathrm{DkC}}(G, k)}{1000 \beta \log ^{3} N}
$$

We obtain a final solution $\mathcal{S}^{*}$ to instance $\operatorname{DkC}(G, k)$ of the Dense $k$-Coloring problem by starting with $\mathcal{S}^{*}=\mathcal{S}^{\prime}$, and then partitioning the vertices of $V(G) \backslash\left(\bigcup_{j=1}^{N / k} S_{j}^{\prime}\right)$ among the sets of $\mathcal{S}^{*}$ arbitrarily, until each such set contains exactly $k$ vertices. Clearly, the value of the resulting solution is at least $\left|E^{\prime \prime}\right|$.

So far we have obtained a randomized algorithm that either returns FAIL (with probability at most $\left.2 / N^{3}\right)$, or it returns a solution to instance $\operatorname{DkC}(G, k)$ of the Dense $k$-Coloring problem, whose expected value is at least $\frac{\operatorname{OPT}_{\mathrm{DkC}}(G, k)}{1000 \beta \log ^{3} N}$.

Let $p^{\prime}$ be the probability that the algorithm returned a solution of value at least $\frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{3000 \beta \log ^{3} N}$, given that it did not return FAIL. Note that the expected solution value, assuming the algorithm did not return FAIL, is at most $\frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{2000 \beta \log ^{3} N}+p^{\prime} \cdot \mathrm{OPT}_{\mathrm{DkC}}(G, k)$. Since this expectation is also at least $\frac{\mathrm{OPT}_{\mathrm{DKC}}(G, k)}{1000 \beta \log ^{3} N}$, we get that $p^{\prime} \geq \frac{1}{2000 \beta \log ^{3} N}$. Overall, the probability that our algorithm successfully returns a solution of value at least $\frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{2000 \beta \log ^{3} N}$ is $p^{\prime} \cdot \operatorname{Pr}[\neg \mathcal{E}] \geq \Omega\left(\frac{1}{\beta \log ^{3} N}\right)$. By repeating the algorithm $\operatorname{poly}(N)$ times we can ensure that it successfully computes a solution of value at least $\frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{2000 \beta \log ^{3} N}$ with high probability.

### 4.2 Approximately Solving the LP-Relaxation

As observed already, (LP-P) has $\binom{N}{k}$ variables, and so we cannot solve it directly. Instead, we will use the Ellipsoids algorithm with an approximate separation oracle to its dual LP, that appears below. This LP has a variable $z$, and, additionally, for every vertex $v \in V(G)$, it has a variable $y_{v}$.

$$
\begin{array}{rcl}
(\mathrm{LP-D}) & & \\
\min & \frac{N}{k} \cdot z+\sum_{v \in V(G)} y_{v} & \\
\text { s.t. } & \\
& z+\sum_{v \in S} y_{v} \geq m(S) & \forall S \in \mathcal{R} \\
& z \geq 0 & \\
& y_{v} \geq 0 & \forall v \in V(G)
\end{array}
$$

We denote the value of the optimal solution to (LP-D) by OPT $_{\text {LP-D }}$.
Next, we define an approximate separation oracle, and provide such a separation oracle to (LP-D).

Approximate Separation Oracle. Consider the following general minimization Linear Program, whose variables are $\left\{x_{1}, \ldots, x_{n}\right\}$.

$$
\begin{array}{ccc}
(P) & & \\
\min & \sum_{i=1}^{n} c_{i} x_{i} & \\
\text { s.t. } & & \\
& \sum_{i=1}^{n} A_{j, i} x_{i} \geq b_{j} & \forall 1 \leq j \leq m \\
& x_{i} \geq 0 & \forall 1 \leq i \leq n
\end{array}
$$

Let $\beta: \mathbb{Z}^{+} \rightarrow \mathbb{Z}^{+}$be an increasing function. A randomized $\beta(n)$-approximate separation oracle for $(\mathrm{P})$ is an efficient randomized algorithm (that is, the running time of the algorithm is bounded by a polynomial function of its input size). The input to the algorithm is a set $\left\{x_{i}\right\}_{i=1}^{n}$ of non-negative real values. The algorithm either returns "accept", or it returns an LP-constraint (called a violated constraint) that does not hold for the given values $x_{1}, \ldots, x_{n}$. We say that the algorithm errs if it returns "accept" and yet there is some index $1 \leq j \leq m$ for which $\sum_{i=1}^{n} A_{j, i} x_{i}<b_{j} / \beta(n)$ holds. We require that the probability that the algorithm errs is at most $2 / 3$.

For the case where a linear program has a very large number of constraints, or the constraints are not given explicitly, one can use an approximate separation oracle, combined with the Ellipsoids algorithm, in order to compute an approximate LP-solution in time polynomial in the number of variables of the LP, provided that there is an Ellipsoid containing the feasible region, whose volume is not too large. For the case where a linear program has a large number of variables, but a relatively small number of constraints, we can use an approximate separation oracle for its dual LP in order to solve the original LP approximately. We start by providing an approximate separation oracle for (LP-D). We then show that this separation oracle can be used in order to obtain an approximate solution to (LP-P) in time poly $(N)$.

Lemma 4.4. Assume that there is an efficient $\alpha(n)$-approximation algorithm for the Densest $k$ Subgraph problem, where $\alpha$ is an increasing function, and $n$ is the number of vertices in the input graph. Then there is a randomized $\beta(N)$-approximate separation oracle for ( $L P-D$ ), where $N$ is the number of variables in the input graph $G$, and $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$.

Proof: Recall that we are given as input real values $z$ and $\left\{y_{v} \mid v \in V(G)\right\}$. Clearly, we can efficiently check whether $z \geq 0$, and whether $y_{v} \geq 0$ for all $v \in V(G)$. If this is not the case, we can return the corresponding violated constraint.
We say that a set $S \in \mathcal{R}$ of vertices is bad if $z+\sum_{v \in S} y_{v}<m(S) / \beta$ holds, where $\beta=c \cdot \alpha\left(N^{2}\right) \cdot \log ^{2} N$, and $c$ is a large enough constant whose value we set later. Our goal is to design an efficient algorithm that either returns a violated constraint of the LP (that is, a set $S \in \mathcal{R}$ of vertices for which $z+\sum_{v \in S} y_{v}<m(S)$ holds); or it returns "accept". We require that, if there exists a bad set $S$ of vertices, then the probability that the algorithm returns "accept" is at most $2 / 3$.
It will be convenient for us to slightly modify the input values in $\left\{y_{v} \mid v \in V(G)\right\}$, as follows. We let $m$ be the smallest integral power of 2 that is greater than $|E(G)|$. First, for every vertex $v \in V(G)$ with $y_{v}>m$, we let $y_{v}^{\prime}=m$, and for every vertex $v \in V(G)$ with $y_{v}<1 / 4$, we set $y_{v}^{\prime}=0$. For each remaining vertex $v$, we let $y_{v}^{\prime}$ be the smallest integral power of 2 that is greater than $4 y_{v}$. Note that for every vertex $v$ with $y_{v}^{\prime} \neq 0,1 \leq y_{v}^{\prime} \leq 4 m$ holds, and $y_{v}^{\prime}$ is an integral power of 2 . We also set $z^{\prime}=2 z$. We say that a set $S \in \mathcal{R}$ of vertices is problematic if $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<8 m(S) / \beta$ holds. We need the following two observations regarding the new values $\left\{y_{v}^{\prime} \mid v \in V(G)\right\}$.

Observation 4.5. If $S \in \mathcal{R}$ is a bad set of vertices, then it is a problematic set of vertices.

Proof: Recall that, if $S$ is a bad set of vertices, then $z+\sum_{v \in S} y_{v}<m(S) / \beta$ must hold. Since, for every vertex $v \in V(G), y_{v}^{\prime} \leq 8 y_{v}$ holds, and $z^{\prime}=2 z$, we get that:

$$
z^{\prime}+\sum_{v \in S} y_{v}^{\prime} \leq 2 z+8 \sum_{v \in S} y_{v} \leq 8\left(z+\sum_{v \in S} y_{v}\right)<8 m(S) / \beta
$$

Therefore, set $S$ is problematic.
Observation 4.6. Assume that there exists a set $S \in \mathcal{R}$ of vertices, for which $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<m(S)$ holds. Let $S^{\prime} \subseteq S$ be the set of vertices obtained from $S$ by deleting every vertex $v \in S$ that has no neighbors in $S$. Then $z+\sum_{v \in S^{\prime}} y_{v}<m\left(S^{\prime}\right)$ holds.

Proof: Since, for every vertex $v \in S \backslash S^{\prime}$, no neighbor of $v$ lies in $S$, we get that $m\left(S^{\prime}\right)=\left|E_{G}\left(S^{\prime}\right)\right|=$ $\left|E_{G}(S)\right|=m(S)$. We partition the vertices of $S^{\prime}$ into two subsets: set $X$ containing all vertices $v \in S^{\prime}$ with $y_{v}<1 / 4$, and set $Y$ containing all remaining vertices. Clearly, $\sum_{v \in X} y_{v}<\frac{|X|}{4} \leq \frac{m\left(S^{\prime}\right)}{2}$ (since $m\left(S^{\prime}\right) \geq\left|S^{\prime}\right| / 2 \geq|X| / 2$, as graph $G\left[S^{\prime}\right]$ contains no isolated vertices).
Assume for contradiction that $z+\sum_{v \in S^{\prime}} y_{v} \geq m\left(S^{\prime}\right)$. Then:

$$
z+\sum_{v \in Y} y_{v} \geq m\left(S^{\prime}\right)-\sum_{v \in X} y_{v} \geq m\left(S^{\prime}\right) / 2
$$

We now consider two cases. The first case is when there is some vertex $v \in Y$ with $y_{v} \geq m$. In this case, $y_{v}^{\prime} \geq m$ holds, and $z^{\prime}+\sum_{v \in S} y_{v}^{\prime} \geq m>m(S)$ holds, a contradiction.
Otherwise, for every vertex $v \in Y, y_{v}^{\prime} \geq 4 y_{v}$ holds. Since $z^{\prime}=2 z$ also holds, we get that:

$$
z^{\prime}+\sum_{v \in S} y_{v}^{\prime} \geq z^{\prime}+\sum_{v \in Y} y_{v}^{\prime} \geq 2 z+4 \sum_{v \in Y} y_{v} \geq m\left(S^{\prime}\right)=m(S)
$$

a contradiction.
From now on we focus on values $z^{\prime},\left\{y_{v}^{\prime} \mid v \in V(G)\right\}$. It is now enough to design an efficient randomized algorithm, that either computes a set $S \in \mathcal{R}$ of vertices, for which $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<m(S)$ holds, or returns "accept". It is enough to ensure that, if there is a problematic set $S \in \mathcal{R}$ of vertices, then the algorithm returns "accept" with probability at most $2 / 3$. Indeed, if there is a bad set $S \in \mathcal{R}$ of vertices, then, from Observation 4.5, there is a problematic set of vertices, and the algorithm will return "accept" with probability at most $2 / 3$. On the other hand, if the algorithm computes a set $S \in \mathcal{R}$ of vertices, for which $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<m(S)$ holds, then we can return the set $S^{\prime} \subseteq S$ of vertices from the statement of Observation 4.6, that defines a violated constraint with respect to the original LP-values.
Our algorithm computes a random partition $(A, B)$ of the vertices of $G$, where every vertex $v \in$ $V(G)$ is independently added to $A$ or to $B$ with probability $1 / 2$ each. Let $q=\log (8 m)$. For all $1 \leq i \leq q$, we define a set $A_{i} \subseteq A$ of vertices: $A_{i}=\left\{v \in A \mid y_{v}^{\prime}=2^{i-1}\right\}$, and we let $A_{0}=\left\{v \in A \mid y_{v}^{\prime}=0\right\}$. Clearly, $\left(A_{0}, \ldots, A_{q}\right)$ is a partition of the set $A$ of vertices.
We compute a partition $\left(B_{0}, \ldots, B_{q}\right)$ of the vertices of $B$ similarly. For all $0 \leq i, j \leq q$, we denote by $E_{i, j}$ the set of all edges $e=(u, v)$ with $u \in A_{i}$ and $v \in B_{j}$, and we define a bipartite graph $G_{i, j}$, whose vertex set is $A_{i} \cup B_{j}$, and edge set is $E_{i, j}$.
Recall that we have assumed that there is an efficient $\alpha(n)$-approximation algorithm for the Densest $k$-Subgraph problem, where $n$ is the number of vertices in the input graph. From Lemma 2.1, there exists an efficient $O\left(\alpha\left(\hat{n}^{2}\right)\right.$ )-approximation algorithm for the Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph problem, where $\hat{n}$ is the number of vertices in the input graph. We denote this algorithm by $\mathcal{A}^{\prime}$.
For every pair $0 \leq i, j \leq q$ of integers, and every pair $k_{1}, k_{2} \geq 0$ of integers with $k_{1}+k_{2} \leq k$, we apply Algorithm $\mathcal{A}^{\prime}$ for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem to graph $G_{i, j}$, with parameters $k_{1}$ and $k_{2}$. Let $S_{i, j}^{k_{1}, k_{2}}$ be the output of this algorithm, and let $m_{i, j}^{k_{1}, k_{2}}$ be the number of edges in the subgraph of $G_{i, j}$ that is induced by the set $S_{i, j}^{k_{1}, k_{2}}$ of vertices. We say that the application of algorithm $\mathcal{A}^{\prime}$ is successful if $z^{\prime}+\sum_{v \in S_{i, j}^{k_{1}, k_{2}}} y_{v}^{\prime}<m_{i, j}^{k_{1}, k_{2}}$, and otherwise it is unsuccessful. If, for any quadruple $\left(i, j, k_{1}, k_{2}\right)$ of indices, the application of algorithm $\mathcal{A}^{\prime}$ was successful, then we return the resulting set $S=S_{i, j}^{k_{1}, k_{2}}$ of vertices. Clearly, $|S| \leq k_{1}+k_{2} \leq k$, so $S \in \mathcal{R}$ holds. Moreover, we are guaranteed that $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<m_{i, j}^{k_{1}, k_{2}}<m(S)$, as required. If every application of algorithm $\mathcal{A}^{\prime}$ is unsuccessful, then we return "accept". The following observation will finish the proof of Lemma 4.4.

Observation 4.7. Suppose there is a problematic set $S \in \mathcal{R}$ of vertices. Then the probability that the algorithm returns "accept" is at most $2 / 3$.

Proof: Let $S \in \mathcal{R}$ be a problematic set of vertices, so $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<8 m(S) / \beta$ holds. Let $E^{\prime}=E_{G}[S]$, so $\left|E^{\prime}\right|=m(S)$. Denote $A_{S}=A \cap S, B_{S}=B \cap S$, and let $E^{\prime \prime} \subseteq E^{\prime}$ be the set of edges $e$, such that exactly one endpoint of $e$ lies in $A$. Clearly, for every edge $e \in E^{\prime}, \operatorname{Pr}\left[e \in E^{\prime \prime}\right]=1 / 2$.

Therefore, $\mathbf{E}\left[\left|E^{\prime \prime}\right|\right]=\left|E^{\prime}\right| / 2$. Let $\mathcal{E}^{\prime}$ be the bad event that $\left|E^{\prime \prime}\right|<\left|E^{\prime}\right| / 8$, and let $p=\operatorname{Pr}\left[\mathcal{E}^{\prime}\right]$. Clearly:

$$
\mathbf{E}\left[\left|E^{\prime \prime}\right|\right] \leq p \cdot \frac{\left|E^{\prime}\right|}{8}+(1-p) \cdot\left|E^{\prime}\right|=\left|E^{\prime}\right|\left(1-\frac{7 p}{8}\right)
$$

Since $\mathbf{E}\left[\left|E^{\prime \prime}\right|\right]=\left|E^{\prime}\right| / 2$, we get that $p \leq 2 / 3$. Next, we show that, if Event $\mathcal{E}^{\prime}$ does not happen, then the algorithm does not return "accept".

From now on we assume that Event $\mathcal{E}^{\prime}$ did not happen, so $\left|E^{\prime \prime}\right| \geq\left|E^{\prime}\right| / 8$. Therefore:

$$
z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<\frac{8 m(S)}{\beta} \leq \frac{64\left|E^{\prime \prime}\right|}{\beta}
$$

holds.
Clearly, there must be a pair $0 \leq i, j \leq q$ of indices, such that $\left|E^{\prime \prime} \cap E_{i, j}\right| \geq \frac{\left|E^{\prime \prime}\right|}{4 q^{2}} \geq \frac{\left|E^{\prime \prime}\right|}{128 \log ^{2} m}$. We now fix this pair $i, j$ of indices, and denote $A_{i}^{\prime}=A_{i} \cap S$ and $B_{j}^{\prime}=B_{j} \cap S$. We also denote $k_{1}=\left|A_{i}^{\prime}\right|$ and let $k_{2}=\left|B_{j}^{\prime}\right|$. Clearly, $k_{1}+k_{2} \leq k$ holds. Denote $M_{i, j}=\left|E^{\prime \prime} \cap E_{i, j}\right|$. From our choice of indices $i, j$, we get that:

$$
z^{\prime}+\sum_{v \in A_{i}^{\prime} \cup B_{j}^{\prime}} y_{v}^{\prime} \leq z^{\prime}+\sum_{v \in S} y_{v}^{\prime} \leq \frac{64\left|E^{\prime \prime}\right|}{\beta} \leq \frac{M_{i, j}}{\beta} \cdot\left(2^{13} \log ^{2} m\right)
$$

Notice that the set $S^{\prime}=A_{i}^{\prime} \cup B_{j}^{\prime}$ of vertices provides a solution to the instance of the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem on graph $G_{i, j}$ with parameters $k_{1}, k_{2}$, whose value is at least $M_{i, j}$. Let $S^{\prime \prime}=S_{i, j}^{k_{1}, k_{2}}$ be the set of vertices obtained by applying Algorithm $\mathcal{A}^{\prime}$ to graph $G_{i, j}$ with parameters $k_{1}, k_{2}$. Since $\left|V\left(G_{i, j}\right)\right| \leq N$, and since $\mathcal{A}^{\prime}$ is an $O\left(\alpha\left(N^{2}\right)\right)$-approximation algorithm for Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph, we are guaranteed that $\left|E_{G}\left(S^{\prime \prime}\right)\right| \geq \Omega\left(\frac{M_{i, j}}{\alpha\left(N^{2}\right)}\right)$. Recall that $\left|A \cap S^{\prime \prime}\right| \leq k_{1}$; $A \cap S^{\prime \prime} \subseteq A_{i}$, and all vertices $v \in A_{i}$ have an identical value $y_{v}^{\prime}$. Therefore, $\sum_{v \in A \cap S^{\prime \prime}} y_{v}^{\prime} \leq \sum_{v \in A_{i}^{\prime}} y_{v}^{\prime}$. Using a similar reasoning, $\sum_{v \in B \cap S^{\prime \prime}} y_{v}^{\prime} \leq \sum_{v \in B_{j}^{\prime}} y_{v}^{\prime}$. Overall, we then get that:

$$
\begin{aligned}
z^{\prime}+\sum_{v \in S^{\prime \prime}} y_{v}^{\prime} & \leq z^{\prime}+\sum_{v \in S^{\prime}} y_{v}^{\prime} \\
& \leq \frac{M_{i, j}}{\beta} \cdot\left(2^{13} \cdot \log ^{2} m\right) \\
& \leq O\left(\frac{\left|E_{G}\left(S^{\prime \prime}\right)\right| \cdot \alpha\left(N^{2}\right) \cdot 2^{13} \cdot \log ^{2} m}{\beta}\right)
\end{aligned}
$$

Recall that $\beta=c \cdot \alpha\left(N^{2}\right) \cdot \log ^{2} N$. By letting the value of the constant $c$ be high enough, we can ensure that $z^{\prime}+\sum_{v \in S^{\prime \prime}} y_{v}^{\prime}<\left|E_{G}\left(S^{\prime \prime}\right)\right|$, and so the application of algorithm $\mathcal{A}^{\prime}$ to graph $G_{i, j}$ with parameters $k_{1}$ and $k_{2}$ is guaranteed to be successful. Therefore, if Event $\mathcal{E}^{\prime}$ does not happen, and we set $c$ to be a large enough constant, then our algorithm does not return "accept". Since $\operatorname{Pr}\left[\mathcal{E}^{\prime}\right] \leq 2 / 3$, the observation follows.

Approximately Solving (LP-P). We use standard methods for solving (LP-P) using approximate separation oracle for (LP-D).

For a collection $\mathcal{R}^{\prime} \subseteq \mathcal{R}$ of vertex subsets, we define a linear program $\left(\operatorname{LP}\left(\mathcal{R}^{\prime}\right)\right.$ ), which is obtained from (LP-D) by only including the constraints associated with the subsets in $\mathcal{R}^{\prime}$ :

$$
\begin{array}{rcl}
\left(\mathrm{LP}\left(\mathcal{R}^{\prime}\right)\right) \\
\min & \frac{N}{k} \cdot z+\sum_{v \in V(G)} y_{v} & \\
\text { s.t. } & \\
& z+\sum_{v \in S} y_{v} \geq m(S) & \forall S \in \mathcal{R}^{\prime} \\
& z \geq 0 & \\
& y_{v} \geq 0 & \forall v \in V(G)
\end{array}
$$

We denote by $\operatorname{OPT}\left(\mathcal{R}^{\prime}\right)$ the value of the optimal solution to $\left(\operatorname{LP}\left(\mathcal{R}^{\prime}\right)\right)$. Since a solution to (LP-D) defines a solution to $\left(\operatorname{LP}\left(\mathcal{R}^{\prime}\right)\right)$, we get that $\operatorname{OPT}\left(\mathcal{R}^{\prime}\right) \leq \mathrm{OPT}_{\mathrm{LP-D}}$ We use the following lemma that allows us to compute a small collection $\mathcal{R}^{\prime} \subseteq \mathcal{R}$ of vertex subsets, such that $\operatorname{OPT}\left(\mathcal{R}^{\prime}\right)$ is within a factor $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$ of $\mathrm{OPT}_{\mathrm{LP}-\mathrm{D}}$.

Claim 4.8. Assume that there is an efficient $\alpha(n)$-approximation algorithm for Densest $k$-Subgraph, where $\alpha$ is an increasing function of $n$, and $n$ is the number of vertices in the input graph. Then there is a randomized algorithm with running time $O(\operatorname{poly}(N))$, that computes a collection $\mathcal{R}^{\prime} \subseteq$ $\mathcal{R}$ of subsets of vertices with $\left|\mathcal{R}^{\prime}\right| \leq O(\operatorname{poly}(N))$, such that, with high probability, $\operatorname{OPT}\left(\mathcal{R}^{\prime}\right) \geq$ $\Omega\left(O P T_{\mathrm{LP}-\mathrm{D}} / \beta(N)\right)$, where $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$.

We prove Claim 4.8 below, after we provide an algorithm for approximately solving (LP-P) using it.

Recall that for every set $S \in \mathcal{R}$ of vertices, there is a variable $x_{S}$ in the primal LP, (LP-P). We now consider the dual $\operatorname{LP}$ to $\operatorname{LP}\left(\mathcal{R}^{\prime}\right)$, that is defined as follows:

$$
\begin{array}{rcl}
\text { (LP-P2) } & & \\
\max & \sum_{S \in \mathcal{R}^{\prime}} m(S) \cdot x_{S} & \\
\text { s.t. } & & \forall v \in V(G) \\
& \sum_{\substack{S \in \mathcal{R}^{\prime} \\
v \in S}} x_{S} \leq 1 & \\
& \sum_{S \in \mathcal{R}^{\prime}} x_{S} \leq N / k \\
& x_{S} \geq 0 & \forall S \in \mathcal{R}^{\prime}
\end{array}
$$

Notice that this Linear Program can be obtained from (LP-P) by eliminating all variables $x_{S}$ for $S \in \mathcal{R} \backslash \mathcal{R}^{\prime}$. Since $\left|\mathcal{R}^{\prime}\right| \leq \operatorname{poly}(N)$, this new linear program has at most $O(\operatorname{poly}(N))$ variables, and it has at most $\operatorname{poly}(N)$ constraints. Therefore, we can solve it in time $O(\operatorname{poly}(N))$ using standard algorithms for LP-solving. Let $\left\{x_{S}^{\prime} \mid S \in \mathcal{R}^{\prime}\right\}$ be the resulting solution. From the Strong Duality Theorem, we get that:

$$
\sum_{S \in \mathcal{R}^{\prime}} m(S) \cdot x_{S}^{\prime}=\operatorname{OPT}\left(\mathcal{R}^{\prime}\right)
$$

From Claim 4.8, with high probability:

$$
\operatorname{OPT}\left(\mathcal{R}^{\prime}\right) \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{LP}-\mathrm{D}}}{\beta(N)}\right)=\Omega\left(\frac{\mathrm{OPT}_{\mathrm{LP-P}}}{\beta(N)}\right)
$$

Altogether, we get that with high probability, $\sum_{S \in \mathcal{R}^{\prime}} m(S) \cdot x_{S}^{\prime} \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{LP}-\mathrm{P}}}{\beta(N)}\right)$. We can extend the solution $\left\{x_{S}^{\prime} \mid S \in \mathcal{R}^{\prime}\right\}$ to (LP-P2) to obtain a feasible solution $\left\{x_{S}^{\prime} \mid S \in \mathcal{R}\right\}$ to (LP-P) by setting the value $x_{S}^{\prime}$ for all sets $S \in \mathcal{R} \backslash \mathcal{R}^{\prime}$ to 0 . It is immediate to verify that the resulting solution to (LP-P) is feasible, and its value remains unchanged. Therefore, we have obtained a randomized algorithm, with running time bounded by $O(\operatorname{poly}(N))$, that with high probability computes a solution to (LP-P), whose value is at least $\Omega\left(\frac{\mathrm{OPT}_{\text {LP-P }}}{\beta(N)}\right)$; here, $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$.
We are now ready to complete the reduction from Dense $k$-Coloring to Densest $k$-Subgraph from Theorem 4.1. Let $\alpha: \mathbb{Z}^{+} \rightarrow \mathbb{Z}^{+}$be an increasing function, such that $\alpha(n) \leq o(n)$, and assume that there is an efficient $\alpha(n)$-approximation algorithm for the Densest $k$-Subgraph problem, where $n$ is the number of vertices in the input graph.

Consider now the input instance $\operatorname{DkC}(G, k)$ of the Dense $k$-Coloring problem with $N=|V(G)|$. We can assume w.l.o.g. that $N$ is greater than a large enough constant, as otherwise we can solve the problem exactly via exhaustive search. Since $\alpha(n) \leq o(n)$, we can also assume that $\beta(N) \leq N^{3}$. We use the randomized algorithm described above, that, in time $O(\operatorname{poly}(N))$, with high probability computes a $\beta(N)$-approximate solution to (LP-P). Recall that the number of variables of (LP-P) with non-zero LP-value is bounded by $O(\operatorname{poly}(N))$. Next, we apply the algorithm from Claim 4.2 in order to round the resulting LP solution. The algorithm with high probability returns an integral solution $\left(S_{1}, \ldots, S_{N / k}\right)$ to instance $\operatorname{DkC}(G, k)$, such that $\sum_{i=1}^{N / k}\left|E_{G}\left(S_{i}\right)\right| \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{DkC}}(G, k)}{\beta(N) \cdot \log ^{3} N}\right)$. Since $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$, with high probability we obtain an $O\left(\alpha\left(N^{2}\right) \cdot\right.$ poly $\left.\log N\right)$-approximate solution to the input instance of Dense $k$-Coloring. In order to complete the reduction from Dense $k$-Coloring to Densest $k$-Subgraph from Theorem 4.1, it is now enough to prove Claim 4.8, which we do next. The proof uses standard techniques and is only included for completeness.
Proof of Claim 4.8. Let $m=|E(G)|$. Notice that we can assume w.l.o.g. that in an optimal solution to (LP-D), for every vertex $v \in V(G), y_{v} \leq m$ holds, and $z \leq m$. Indeed, if this is not the case, then we can modify the solution by setting $y_{v}=\min \left\{y_{v}, m\right\}$ for every vertex $v \in V(G)$, and $z=\min \{z, m\}$. It is easy to verify that this remains a feasible solution, and its value does not grow. For convenience, for every subset $\mathcal{R}^{\prime} \subseteq \mathcal{R}$ of vertex subsets, we define the following LP:

$$
\begin{array}{rcl}
\left(P^{\prime}\left(\mathcal{R}^{\prime}\right)\right) & & \\
\min & \frac{N}{k} \cdot z+\sum_{v \in V(G)} y_{v} & \\
\text { s.t. } & & \\
& z+\sum_{v \in S} y_{v} \geq m(S) & \forall S \in \mathcal{R}^{\prime} \\
& 0 \leq z \leq m & \\
& 0 \leq y_{v} \leq m & \forall v \in V(G)
\end{array}
$$

We denote by $\operatorname{OPT}^{\prime}\left(\mathcal{R}^{\prime}\right)$ the value of the optimal solution of the above LP. From the above discussion, it is enough to provide a randomized algorithm with running time $O(\operatorname{poly}(N))$, that computes a collection $\mathcal{R}^{\prime} \subseteq \mathcal{R}$ of subsets of vertices, such that, with high probability, $\mathrm{OPT}^{\prime}\left(\mathcal{R}^{\prime}\right) \geq$ $\Omega\left(\mathrm{OPT}^{\prime}(\mathcal{R}) / \beta(N)\right)$, where $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$. This is since $\operatorname{OPT}\left(\mathcal{R}^{\prime}\right)=\operatorname{OPT}^{\prime}\left(\mathcal{R}^{\prime}\right)$, and $\mathrm{OPT}^{\prime}(\mathcal{R})=\mathrm{OPT}(\mathcal{R})$ holds.

Clearly, $\operatorname{OPT}^{\prime}(\mathcal{R}) \leq 2 N m$ must hold. Let $M$ be the smallest integral power of 2 that is greater than $2 N m$. For all $0<C \leq \log M$ and $\mathcal{R}^{\prime} \subseteq \mathcal{R}$, we consider a feasibility LP, that is obtained from $P^{\prime}\left(\mathcal{R}^{\prime}\right)$, by adding the constraint that $\frac{N}{k} \cdot z+\sum_{v \in V(G)} y_{v} \leq 2^{C}$ :

$$
\begin{aligned}
& \left(F\left(\mathcal{R}^{\prime}, C\right)\right) \\
& \qquad \begin{array}{cl}
\frac{N}{k} \cdot z+\sum_{v \in V(G)} y_{v} \leq 2^{C} & \\
z+\sum_{v \in S} y_{v} \geq m(S) & \forall S \in \mathcal{R}^{\prime} \\
0 \leq z \leq m & \\
0 \leq y_{v} \leq m & \forall v \in V(G)
\end{array}
\end{aligned}
$$

The key to the proof of Claim 4.8 is the following observation.
Observation 4.9. Assume that there is an efficient $\alpha(n)$-approximation algorithm for Densest $k$-Subgraph, where $\alpha$ is an increasing function of $n$, and $n$ is the number of vertices in the input graph. Then there is a randomized algorithm with running time $O(\operatorname{poly}(N))$, that, given a value $0 \leq C \leq \log M$, either:

- computes a collection $\mathcal{R}^{\prime}(C) \subseteq \mathcal{R}$ of at most $O(\operatorname{poly}(N))$ subsets of vertices, such that the linear program $\left(F\left(\mathcal{R}^{\prime}, C\right)\right)$ is infeasible; or
- computes values $0 \leq z^{\prime} \leq m$ and $0 \leq y_{v}^{\prime} \leq m$ for all $v \in V(G)$, such that $\frac{N}{k} \cdot z+\sum_{v \in V(G)} y_{v} \leq$ $2^{C}$, and, with high probability, for every vertex set $S \in \mathcal{R}, z^{\prime}+\sum_{v \in S} y_{v}^{\prime} \geq m(S) / \beta(N)$.

We provide the proof of Observation 4.9 below, after we complete the proof of Claim 4.8 using it. We apply the algorithm from Observation 4.9 to every value $0 \leq C \leq \log M$. We say that the application of the algorithm for value $C$ is successful if the algorithm returns values $0 \leq z^{\prime} \leq m$ and $0 \leq y_{v}^{\prime} \leq m$ for all $v \in V(G)$; otherwise we say that it is unsuccessful. We say that the algorithm errs if it is successful, and yet there is a set $S \in \mathcal{R}$ of vertices for which $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<m(S) / \beta(N)$.

Let $C^{*}$ be the smallest value of $C$, such that the algorithm from Observation 4.9, when applied to $C$, was successful. Let $z^{\prime}$ and $\left\{y_{v}^{\prime} \mid v \in V(G)\right\}$ be the values of the LP-variables returned by the algorithm. We let $\hat{\mathcal{E}}$ be the bad event that the algorithm from Observation 4.9, when applied to value $C^{*}$, errs. The probability of $\hat{\mathcal{E}}$ is at most $1 / \operatorname{poly}(N)$. Consider the following solution to $\left(P^{\prime}(\mathcal{R})\right)$ : we set $z^{*}=\min \left\{z^{\prime} \cdot \beta(N), m\right\}$, and for all $v \in V(G)$, we set $y_{v}^{*}=\min \left\{y_{v}^{\prime} \cdot \beta(N), m\right\}$. It is immediate to verify that, if Event $\hat{\mathcal{E}}$ did not happen, then we obtain a feasible solution to $\left(P^{\prime}(\mathcal{R})\right)$, whose value is at most $2^{C^{*}} \cdot \beta(N)$. Notice that, from the choice of the value $C^{*}, \operatorname{LP}\left(F\left(\mathcal{R}^{\prime}\left(C^{*}-1\right), C^{*}-1\right)\right)$ does not have a feasible solution. Since the constraints in $\left(F\left(\mathcal{R}^{\prime}\left(C^{*}-1\right), C^{*}-1\right)\right.$ are a subset of the constraints in $\left(F\left(\mathcal{R}, C^{*}-1\right)\right)$, it follows that $\left(F\left(\mathcal{R}, C^{*}-1\right)\right)$ does not have a feasible solution, and so $\operatorname{OPT}^{\prime}(\mathcal{R}) \geq 2^{C^{*}-1}$. Therefore, if Event $\hat{\mathcal{E}}$ did not happen, we obtain a feasible solution $\left(z^{*},\left\{y_{v}^{*}\right\}_{v \in V(G)}\right)$ to $\left(P^{\prime}(\mathcal{R})\right)$, whose value is at most $2 \beta(N) \cdot \mathrm{OPT}^{\prime}(\mathcal{R})$.
The final collection of subsets of vertices that our algorithm returns is $\mathcal{R}^{\prime}=\bigcup_{C=0}^{C^{*}-1} \mathcal{R}^{\prime}(C)$. Clearly, $\left|\mathcal{R}^{\prime}\right| \leq \operatorname{poly}(N)$. It is also immediate to verify that $\operatorname{OPT}\left(\mathcal{R}^{\prime}\right) \geq 2^{C^{*}-1}$. Indeed, for all values $0 \leq C<C^{*}$, linear program $\left(F\left(\mathcal{R}^{\prime}(C), C\right)\right)$ is infeasible, and, since $\mathcal{R}^{\prime}(C) \subseteq \mathcal{R}^{\prime}$, linear program $\left(F\left(\mathcal{R}^{\prime}, C\right)\right)$ is also infeasible. Since every constraint of $P^{\prime}\left(\mathcal{R}^{\prime}\right)$ is also a constraint of $L P(\mathcal{R})$, we get that, if Event $\hat{\mathcal{E}}$ did not happen, then $\left(z^{*},\left\{y_{v}^{*}\right\}_{v \in V(G)}\right)$ is a feasible solution to $P^{\prime}\left(\mathcal{R}^{\prime}\right)$, whose value is at most $\beta(N) \cdot 2^{C^{*}}$. To summarize, if Event $\hat{\mathcal{E}}$ did not happen, then:

$$
2^{C^{*}-1} \leq \mathrm{OPT}^{\prime}\left(\mathcal{R}^{\prime}\right) \leq 2^{C^{*}} \cdot \beta(N)
$$

and

$$
2^{C^{*}-1} \leq \mathrm{OPT}^{\prime}(\mathcal{R}) \leq 2^{C^{*}} \cdot \beta(N)
$$

hold.
Therefore, if Event $\hat{\mathcal{E}}$ did not happen, then $\operatorname{OPT}^{\prime}\left(\mathcal{R}^{\prime}\right) \geq 2^{C^{*}-1} \geq \mathrm{OPT}^{\prime}(\mathcal{R}) /(2 \beta(N))$. It now remains to prove Observation 4.9. The proof is standard; we only provide its sketch below.
Proof of Observation 4.9. We fix a value $0 \leq C \leq \log M$, and consider the corresponding Linear Program $(F(\mathcal{R}, C))$. The idea of the proof is simple: we employ the Ellipsoids algorithm, together with the separation oracle from Lemma 4.4 (after we reduce its error probability by repeating the algorithm a number of times). We then let $\mathcal{R}^{\prime}(C)$ be the collection of all vertex subsets $S \in \mathcal{R}$, such that the separation oracle returns the constraint associated with $S$ over the course of the algorithm.
We now provide more details. Recall first the Ellipsoids algorithm for solving a feasibility Linear Program $(F)$ on $N^{\prime}$ variables. The algorithm proceeds in iterations. The input to the $i$ th iteration is an $N^{\prime}$-dimensional Ellipsoid $E_{i}$, that contains the feasible region of $(F)$. Let $x_{i}$ denote the center point of the ellipsoid. If the algorithm is given a constraint $A^{j}$ of the Linear Program $(F)$ that is violated by point $x_{i}$, then it produces a new ellipsoid $E_{i+1}$, that contains the feasible region of $(F)$, whose volume is at most $\left(1-1 / \operatorname{poly}\left(N^{\prime}\right)\right)$ times the volume of $E_{i}$. The running time of a single iteration is $O\left(\operatorname{poly}\left(N^{\prime}\right)\right)$.
Typically, we assume that there is an initial ellipsoid $E_{1}$, whose volume is at most $2^{\text {poly }\left(N^{\prime}\right)}$, that contains the feasible region of $(F)$, which needs to be supplied to the Ellipsoids algorithm. We can also typically assume that, if $(F)$ has a feasible solution, then the volume of the feasible region of $(F)$ is at least $L=2^{-\operatorname{poly}\left(N^{\prime}\right)}$ (if this is not the case, the feasible region can be slightly inflated artificially
by adding a small amount of slack to the constraints; in our case, since we are only solving the LP approximately, this is immaterial). If the above two conditions hold, the algorithm can proceed for at most poly $\left(N^{\prime}\right)$ iterations, before the volume of the current ellipsoid becomes smaller than $L$, and the algorithm then correctly declares that $(F)$ does not have a feasible solution. In every iteration, the constraint violated by the center $x_{i}$ of the current ellipsoid $E_{i}$ is supplied by a separation oracle. If the separation oracle declares that $x_{i}$ is (approximately) feasible solution, then the algorithm halts.

We now turn to consider the linear program $(F(\mathcal{R}, C))$. Since the LP constraints require that the values of all LP-variables are between 0 and $m$, it is easy to verify that the feasible region of the LP is contained in the $(N+1)$-dimensional sphere $E_{1}$, whose radius is bounded by poly $(m)$, and volume is at most $2^{\text {poly }(m)} \leq 2^{\text {poly }(N)}$. We initially set $\mathcal{R}^{\prime}(C)=\emptyset$. We apply the Ellipsoids algorithm to this LP, with the initial ellipsoid $E_{1}$.

We now consider the $i$ th iteration of the algorithm, whose input is an ellipsoid $E_{i}$, together with its center point $\left(z^{i},\left\{y_{v}^{i}\right\}_{v \in V(G)}\right)$. We manually check the constraints $\frac{N}{k} \cdot z^{i}+\sum_{v \in V(G)} y_{v}^{i} \leq 2^{C}$; $0 \leq z^{i} \leq m$; and $0 \leq y_{v}^{i} \leq m$ for all $v \in V(G)$. If any of these constraints does not hold, then we return it as a violated constraint. Assume now that all these constraints hold. We apply the algorithm from Lemma 4.4 to the current values $\left(z^{i},\left\{y_{v}^{i}\right\}_{v \in V(G)}\right)$; we do so $N$ times. If, in each of these iterations, the algorithm returns "accept", then we terminate our algorithm, and return the current solution $\left(z^{i},\left\{y_{v}^{i}\right\}_{v \in V(G)}\right)$. Observe that we are guaranteed that $0 \leq z^{i} \leq m ; 0 \leq y_{v}^{i} \leq m$ for all $v \in V(G)$; and $\frac{N}{k} \cdot z^{i}+\sum_{v \in V(G)} y_{v}^{i} \leq 2^{C}$. Moreover, unless the algorithm from Lemma 4.4 erred in each of its $N$ applications, we are guaranteed that, for every vertex set $S \in \mathcal{R}, z^{i}+\sum_{v \in S} y_{v}^{i} \geq$ $m(S) / \beta(N)$ holds. The probability that the algorithm from Lemma 4.4 errs is at most $2 / 3$, so with high probability, we are guaranteed that for all $S \in \mathcal{R}, z^{i}+\sum_{v \in S} y_{v}^{i} \geq m(S) / \beta(N)$.
Assume now that in some application of the algorithm from Lemma 4.4 to the current values $\left(z^{i},\left\{y_{v}^{i}\right\}_{v \in V(G)}\right)$ we obtain a violated constraint of (LP-D). That is, we obtain a set $S \in \mathcal{R}$ of vertices, for which $z^{i}+\sum_{v \in S} y_{v}^{i}<m(S)$ holds. In this case, we add $S$ to set $\mathcal{R}^{\prime}(S)$, and we use this constraint as a violated constraint for the Ellipsoids algorithm.
If the above algorithm never terminates with an approximately feasible solution $\left(z^{i},\left\{y_{v}^{i}\right\}_{v \in V(G)}\right)$, then we are guaranteed that after at most $\operatorname{poly}(N)$ iterations, the algorithm correctly certifies that $(F(\mathcal{R}, C))$ does not have a feasible solution. We then return the current collection $\mathcal{R}(C)$ of vertex subsets. For convenience, we denote by $A^{1}, A^{1}, \ldots, A^{r}$ the sequence of violated constraints that were fed to the Ellipsoids algorithm. Each of the constraints $A^{j}$ either corresponds to a set $S \in \mathcal{R}(C)$, or it is one of the constraints $0 \leq z^{i} \leq m ; 0 \leq y_{v}^{i} \leq m$ for all $v \in V(G)$; and $\frac{N}{k} \cdot z^{i}+\sum_{v \in V(G)} y_{v}^{i} \leq 2^{C}$. In other words, each such constraint $A^{j}$ is also a constraint of the LP $(F(\mathcal{R}(C), C))$.
It now remains to prove that in the latter case, $(F(\mathcal{R}(C), C))$ does not have a feasible solution. In order to do so, consider applying the Ellipsoids algorithm to this linear program. We start with the same initial ellipsoid $E_{1}$ as before. Since the Ellipsoid algorithm is deterministic, its behavior is entirely determined by the initial ellipsoid $E_{1}$ and the sequence of the violated constraints that it receives. We will use exactly the same sequence $A^{1}, A^{1}, \ldots, A^{r}$ of violated constraints in this execution of Ellipsoids algorithm. This ensures that for all $i$, the ellipsoid $E_{i}$ that is used as the input to the $i$ th iteration is identical to the ellipsoid that was used as input to iteration $i$ when solving $(F(\mathcal{R}, C))$, which in turn ensures that constraint $A^{i}$ is a violating constraint for the center of ellipsoid $E_{i}$. Therefore, this execution of Ellipsoids algorithm is identical to the execution of the
same algorithm when applied to $\operatorname{LP}(F(\mathcal{R}, C))$, and it will end up with a final ellipsoid $E_{r}$, whose volume is small enough to correctly establish that $(F(\mathcal{R}(C), C))$ does not have a feasible solution.

## 5 Reductions from Densest $k$-Subgraph to Dense $k$-Coloring and (r,h)Graph Partitioning

In this section we prove the following theorem.
Theorem 5.1. Let $\alpha: \mathbb{Z}^{+} \rightarrow \mathbb{Z}^{+}$be an increasing function with $\alpha(n) \leq o(n)$. Then the following hold:

- If there exists an efficient $\alpha(n)$-approximation algorithm $\mathcal{A}$ for the Dense $k$-Coloring problem, where $n$ is the number of vertices in the input graph, then there exists a randomized algorithm for the Densest $k$-Subgraph problem, whose running time is $N^{(\log N)}$, that with high probability computes an $O\left(\alpha\left(N^{O(\log N)}\right) \cdot \log N\right)$-approximate solution to the input instance of the problem; here $N$ is the number of vertices in the input instance of Densest $k$-Subgraph.
- If there exists an efficient $\alpha(n)$-approximation algorithm for the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning problem, where $n$ is the number of vertices in the input graph, then exists a randomized algorithm for the Densest $k$-Subgraph problem, whose running time is $N^{O(\log N)}$, that with high probability computes an $O\left(\left(\alpha\left(N^{O(\log N)}\right)\right)^{3} \cdot \log ^{2} N\right)$-approximate solution to the input instance of the problem; here $N$ is the number of vertices in the input instance of Densest $k$-Subgraph.

We obtain the following immediate corollary of Theorem 5.1.
Corollary 5.2. Assume that Conjecture 2 holds and that NP $\nsubseteq \operatorname{BPTIME}\left(n^{O(\log n)}\right)$. Then for some constant $0<\varepsilon^{\prime} \leq 1 / 2$, there is no efficient $2^{(\log n)^{\varepsilon^{\prime}}}$-approximation algorithm for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, and there is no efficient $2^{(\log n)^{\varepsilon^{\prime}}}$-approximation algorithm for Dense $k$-Coloring.

Proof: We prove the corollary for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning; the proof for Dense $k$-Coloring is similar. Assume that Conjecture 2 holds and that NP $\nsubseteq \operatorname{DTIME}\left(n^{O(\log n)}\right)$. Then, from Theorem 3.2, for some constant $0<\varepsilon<1$, there is no randomized factor-2 $2^{(\log n)^{\varepsilon}}$-approximation algorithm for Densest $k$-Subgraph with running time $n^{O(\log n)}$, where $n$ is the number of vertices in the input graph.
We let $\varepsilon^{\prime}=\varepsilon / c$, where $c$ is a sufficiently large constant. We now prove that there is no efficient $2^{(\log n)^{\varepsilon^{\prime}}}$-approximation algorithm for $(r, h)$-Graph Partitioning. Indeed, assume for contradiction that there is an efficient $2^{(\log n)^{\varepsilon^{\prime}}}$-approximation algorithm $\mathcal{A}$ for $(r, h)$-Graph Partitioning. From Theorem 5.1, there is a randomized algorithm for the Densest $k$-Subgraph problem, that, given an instance $\operatorname{DkS}(G, k)$ of the problem with $|V(G)|=N$, in time $N^{O(\log N)}$, computes a $c^{\prime}\left(\alpha\left(N^{O(\log N)}\right)\right)^{3} \cdot \log ^{2} N$ approximate solution, where $\alpha(x)=2^{(\log x)^{\varepsilon^{\prime}}}$ and $c^{\prime}$ is a constant independent of $N$. Note that:

$$
\alpha\left(N^{O(\log N)}\right)=2^{(\log N)^{O\left(\varepsilon^{\prime}\right)}}=2^{(\log N)^{O(\varepsilon / c)}} .
$$

Since we can let $c$ be a sufficiently large constant, we can ensure that $c^{\prime}\left(\alpha\left(N^{O(\log N)}\right)\right)^{3} \cdot \log ^{2} N<$ $2^{(\log n)^{\varepsilon}}$.

Therefore, we obtain a randomized factor-2 $2^{(\log n)^{\varepsilon}}$-approximation algorithm for Densest $k$-Subgraph, with running time $n^{O(\log n)}$, a contradiction.

The remainder of this section is dedicated to the proof of Theorem 5.1. In order to obtain both reductions, we start with an instance $\operatorname{DkS}(G, k)$ of the Densest $k$-Subgraph problem, and construct another auxiliary graph $H$. This graph is then used in order to define the corresponding instances of Dense $k$-Coloring and (r,h)-Graph Partitioning, respectively. We start by defining graph $H$ and analyzing its properties in Section 5.1. We then complete the reduction from Densest $k$-Subgraph to Dense $k$-Coloring in Section 5.2, and the reduction from Densest $k$-Subgraph to ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning in Section 5.3. Throughout this section, for an integer $A$, we denote $[A]=\{0,1, \ldots, A-1\}$. We also assume that the parameter $k$ in the input instance of the Densest $k$-Subgraph problem is greater than a large enough constant, since otherwise the problem can be solved in time poly $(N)$ via exhaustive search.

### 5.1 Auxiliary Graph $H$

Let $\operatorname{DkS}(G, k)$ be an instance of the Densest $k$-Subgraph problem. Denote $V(G)=\left\{v_{0}, v_{1}, \ldots, v_{N-1}\right\}$. We now provide a randomized algorithm to construct an auxiliary graph $H$ corresponding to this instance. The construction is somewhat similar to and inspired by the construction used in Section 2 of KLS00].
Let $q=\lceil\log N\rceil$. We start by computing a prime number $M$, such that $N^{5 q} \leq M \leq 2 \cdot N^{5 q}$. From the Bertrand-Chebyshev theorem Ber45, Ceb50, such a prime number must exist, and it can be computed in time $N^{O(\log N)}$ by checking every integer between $N^{5 q}$ and $2 \cdot N^{5 q}$. We then construct a random mapping $f:[N] \rightarrow[M]$ as follows. For every integer $1 \leq i \leq N$, we let $f(i)$ be an integer chosen independently and uniformly at random (with replacement) from $[M]$.
We are now ready to define the graph $H$. The set of vertices of $H$ is $V(H)=\left\{u_{0}, \ldots, u_{M-1}\right\}$. For every edge $e=\left(v_{i}, v_{j}\right) \in E(G)$, we construct a collection $J(e)$ of $M$ edges in $H: J(e)=$ $\left\{\left(u_{f(i)+t}, u_{f(j)+t}\right) \mid 0 \leq t \leq M-1\right\}$, where the addition in the subscript is modulo $M$ (we use this convention throughout the remainder of this section). We say that edge $e$ is the origin of every edge in set $J(e)$. We then set $E(H)=\bigcup_{e \in E(G)} J(e)$. We note that we do not allow parallel edges in $H$, so it is possible for an edge in $H$ to have several origin edges in $G$. This completes the definition of the graph $H$. We now analyze its properties.

Good Event $\mathcal{E}^{g}$. We say that a good event $\mathcal{E}^{g}$ happens if there is a collection $\left\{H_{1}, \ldots, H_{r}\right\}$ of $r=\left\lfloor\frac{M}{k \log k}\right\rfloor$ disjoint subgraphs of $H$, such that the following hold:

- for all $1 \leq j \leq r,\left|V\left(H_{j}\right)\right|=k$;
- for all $1 \leq j \leq r,\left|E\left(H_{j}\right)\right| \leq \mathrm{OPT}_{\mathrm{DkS}}(G, k)$; and
- $\sum_{1 \leq j \leq r}\left|E\left(H_{j}\right)\right| \geq 0.1 \cdot\left\lfloor\frac{M}{k \log k}\right\rfloor \cdot \operatorname{OPT}_{\mathrm{DkS}}(G, k)$.

We start by showing that good event $\mathcal{E}^{g}$ happens with a sufficiently high probability.
Claim 5.3. $\operatorname{Pr}\left[\mathcal{E}^{g}\right] \geq 0.8$.

Proof: Let $S$ be the optimal solution to instance $\operatorname{DkS}(G, k)$ of the Densest $k$-Subgraph problem, and let $T$ be a subset of vertices of $H$, defined as $T=\left\{u_{f(i)} \mid v_{i} \in S\right\}$. Let $I$ be a collection of $r=\lfloor M /(k \log k)\rfloor$ integers from $[M]$, obtained by selecting each integer independently uniformly at random (with replacement) from $[M]$. For every index $j \in I$, we define a set $T_{j}^{\prime}$ of vertices of $H$ as follows: $T_{j}^{\prime}=\left\{u_{f(i)+j} \mid v_{i} \in S\right\}$. Finally, for every index $j \in I$, we define another set $T_{j}$ of vertices of $H$, by starting with the set $T_{j}^{\prime}$ of vertices, and then removing from it every vertex that lies in set $\bigcup_{i \in I \backslash\{j\}} T_{i}^{\prime}$. Clearly, all resulting vertex sets in $\left\{T_{j}\right\}_{j \in I}$ are mutually disjoint, and each such set contains at most $k$ vertices.
For every index $j \in I$, we denote by $\hat{E}\left(T_{j}\right)$ the set of edges $e \in E_{H}\left(T_{j}\right)$, such that an origin of $e$ in $G$ lies in $E_{G}(S)$. Equivalently: $\hat{E}\left(T_{j}\right)=\left\{\left(u_{f(i)+j}, u_{f\left(i^{\prime}\right)+j}\right) \mid\left(v_{i}, v_{i^{\prime}}\right) \in E_{G}(S)\right\}$. Clearly, $\left|\hat{E}\left(T_{j}\right)\right| \leq$ $\mathrm{OPT}_{\mathrm{DkS}}(G, k)$ for all $j \in I$, and so $\sum_{j \in I}\left|\hat{E}\left(T_{j}\right)\right| \leq r \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)=\left\lfloor\frac{M}{k \log k}\right\rfloor \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)$.
We prove the following observation.

## Observation 5.4.

$$
\mathbf{E}\left[\sum_{j \in I}\left|\hat{E}\left(T_{j}\right)\right|\right] \geq 0.9 \cdot\left\lfloor\frac{M}{k \log k}\right\rfloor \cdot \operatorname{OPT}_{\mathrm{DkS}}(G, k) .
$$

Assume first that the observation holds. For simplicity of notation, denote $B=\left\lfloor\frac{M}{k \log k}\right\rfloor \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)$. For all $j \in I$, we define a subgraph $H_{j}$ of $H$, whose vertex set is $T_{j}$, and edge set is $\hat{E}\left(T_{j}\right)$. Clearly, the graphs in $\left\{H_{j} \mid j \in I\right\}$ are disjoint, and, for all $j \in I,\left|V\left(H_{j}\right)\right|=|T| \leq k$ holds. Additionally, from the above discussion, for all $j \in I,\left|E\left(H_{j}\right)\right| \leq \operatorname{OPT}_{\mathrm{DkS}}(G, k)$. If, additionally, $\sum_{j \in I}\left|\hat{E}\left(T_{j}\right)\right| \geq 0.1 B$ holds, then Event $\mathcal{E}^{g}$ happens. As observed above, $\sum_{j \in I}\left|\hat{E}\left(T_{j}\right)\right| \leq B$ must hold. Let $p$ denote the probability that $\sum_{j \in I}\left|\hat{E}\left(T_{j}\right)\right| \geq 0.1 B$. Then:

$$
\mathbf{E}\left[\sum_{j \in I}\left|\hat{E}\left(T_{j}\right)\right|\right] \leq(1-p) \cdot 0.1 \cdot B+p \cdot B=0.1 B+0.9 B p .
$$

Since, from Observation 5.4. $\mathbf{E}\left[\sum_{j \in I}\left|\hat{E}\left(T_{j}\right)\right|\right] \geq 0.9 B$, we conclude that $p \geq 0.8$, and so $\operatorname{Pr}\left[\mathcal{E}^{g}\right] \geq$ 0.8 , as required. In order to complete the proof of Claim 5.3, it is now enough to prove Observation 5.4, which we do next.
Proof of Observation 5.4. We associate a collection $\left\{X_{1}, \ldots, X_{r}\right\}$ of random variables with the set $I$ of indices. In order to do so, we view set $I$ as being constructed as follows. For each $1 \leq i \leq r$, sample a value $X_{i}$ uniformly at random from $[M]$, and then let $I=\left\{X_{1}, \ldots, X_{r}\right\}$ be the collection of these sampled values. Consider now any index $1 \leq i \leq r$, the corresponding set $T_{X_{i}}^{\prime}$ of vertices of $H$, and any edge $e=\left(v_{a}, v_{b}\right) \in E_{G}(S)$. Edge $e^{\prime}=\left(u_{f(a)+X_{i}}, u_{f(b)+X_{i}}\right)$ of $H$ corresponding to $e$ belongs to set $\hat{E}\left(T_{X_{i}}\right)$ if and only if neither of the vertices $u_{f(a)+X_{i}}, u_{f(b)+X_{i}}$ lies in $\bigcup_{j \in I \backslash\{i\}} T_{X_{j}}^{\prime}$. Consider now an index $j \in I \backslash\{i\}$. The probability that a fixed vertex $u \in V(H)$ lies in $T_{X_{j}}^{\prime}$ is at most $k / M$
(since for every vertex $v_{z} \in S$, there is a single index $s$ with $u_{f(z)+s}=u$ ). From the union bound, the probability that a fixed vertex $u \in V(H)$ lies in $\bigcup_{j \in I \backslash\{i\}} T_{X_{j}}^{\prime}$ is at most $\frac{r \cdot k}{M} \leq \frac{1}{\log k}$. In particular, the probability that any of the endpoints of edge $e^{\prime}=\left(u_{f(a)+X_{i}}, u_{f(b)+X_{i}}\right)$ lies in $\bigcup_{j \in I \backslash\{i\}} T_{X_{j}}^{\prime}$ is at most $\frac{2}{\log k}$. Therefore, $\mathbf{E}\left[\left|\hat{E}\left(T_{j}\right)\right|\right] \geq\left|E_{G}(S)\right| \cdot\left(1-\frac{2}{\log k}\right) \geq 0.9 \cdot\left|E_{G}(S)\right|$. Altogether, from the linearity of expectation, $\mathbf{E}\left[\sum_{1 \leq i \leq r}\left|\hat{E}\left(T_{X_{i}}\right)\right|\right] \geq 0.9 \cdot\left|E_{G}(S)\right| \cdot r=0.9 \cdot\left\lfloor\frac{M}{k \log k}\right\rfloor \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)$.

Ensemble and Bad Event $\mathcal{E}^{b}$. Next, we define the notion of an ensemble. Recall that $q=$ $\lceil\log N\rceil$. An ensemble $\mathcal{B}$ consists of a collection $I \subseteq[N]$ of at most $2 q$ indices, and, for every index $i \in I$, an integer $-q \leq x_{i} \leq q$ with $x_{i} \neq 0$. We denote the ensemble by $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$. We say that ensemble $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$ is bad if $\sum_{i \in I} x_{i} \cdot f(i) \equiv 0(\bmod M)$. We let $\mathcal{E}^{b}$ be the bad event that there exists a bad ensemble. We start by showing that the probability of Event $\mathcal{E}^{b}$ happening is low. Later, we show that, if Event $\mathcal{E}^{b}$ does not happen, then graph $H$ has some useful properties.

Observation 5.5. $\operatorname{Pr}\left[\mathcal{E}^{b}\right] \leq \frac{1}{N^{q}}$.

Proof: Consider any fixed ensemble $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$. Let $i^{*} \in I$ be any fixed index. Consider now the following two-step process: in the first step, we select the values $f(i)$ for all indices $i \in I \backslash\left\{i^{*}\right\}$ independently uniformly at random from $[M]$. We then denote $S=\sum_{i \in I \backslash\left\{i^{*}\right\}} x_{i} \cdot f(i)$. In the second step, we select the value $a=f\left(i^{*}\right)$ at random from $[M]$. Ensemble $\mathcal{B}$ is bad if and only if $a \cdot x_{i^{*}}+S=0 \bmod M$. Since $M$ is a prime number, there is exactly one value $a^{\prime} \in[M]$ with $a^{\prime} \cdot x_{i^{*}}+S=0 \bmod M$ (indeed, if two such values $a^{\prime}, a^{\prime \prime}$ exist, then $a^{\prime} \cdot x_{i^{*}}=a^{\prime \prime} \cdot x_{i^{*}} \bmod M$, implying that $a^{\prime}=a^{\prime \prime}$ must hold). The probability to choose $f\left(i^{*}\right)=a^{\prime}$ is then $1 / M$, and so the probability that a fixed ensemble $\mathcal{B}$ is bad is $1 / M$.
Notice that the total number of ensembles is bounded by $\left(\sum_{t=1}^{2 q}\binom{N}{t}\right) \cdot(2 q)^{2 q} \leq(2 q) \cdot\binom{N}{2 q} \cdot(2 q)^{2 q} \leq$ $N^{4 q}$. Using the Union Bound, $\operatorname{Pr}\left[\mathcal{E}^{b}\right] \leq \frac{N^{4 q}}{M} \leq \frac{1}{N^{q}}$, since $M \geq N^{5 q}$.
Next, we show that, if Event $\mathcal{E}^{b}$ does not happen, then every edge $e \in E(H)$ has a unique origin edge in $G$.

Observation 5.6. Assume that Event $\mathcal{E}^{b}$ did not happen. Let $e$ be any edge of $H$. Then there is a unique edge $e^{\prime} \in E(G)$, such that $e^{\prime}$ is an origin edge of $e$.

Proof: Denote $e=\left(u_{j}, u_{j^{\prime}}\right)$, and assume for contradiction that there are two distinct edges $e_{1}, e_{2} \in$ $E(G)$ that both serve as origin edges of $e$. Denote $e_{1}=\left(v_{i_{1}}, v_{i_{1}^{\prime}}\right)$ and $e_{2}=\left(v_{i_{2}}, v_{i_{2}^{\prime}}\right)$. From the construction of $H$, since edge $e_{1}$ is an origin edge of $e$, there exists an integer $t_{1}$, such that:

$$
\begin{equation*}
j \equiv f\left(i_{1}\right)+t_{1}(\bmod M), \quad \text { and } \quad j^{\prime} \equiv f\left(i_{1}^{\prime}\right)+t_{1}(\bmod M) . \tag{5}
\end{equation*}
$$

Similarly, since edge $e_{2}$ is an origin edge of $e$, there exists an integer $t_{2}$, such that:

$$
\begin{equation*}
j \equiv f\left(i_{2}\right)+t_{2}(\bmod M), \quad \text { and } \quad j^{\prime} \equiv f\left(i_{2}^{\prime}\right)+t_{2}(\bmod M) . \tag{6}
\end{equation*}
$$

By adding the first equation of (5) to the second equation of (6), we get that $f\left(i_{1}\right)+f\left(i_{2}^{\prime}\right) \equiv$ $j+j^{\prime}-t_{1}-t_{2}(\bmod M)$. Similarly, by adding the second equation of $(5)$ to the first equation of (6), we get that $f\left(i_{1}^{\prime}\right)+f\left(i_{2}\right) \equiv j+j^{\prime}-t_{1}-t_{2}(\bmod M)$. In other words, we get that:

$$
\begin{equation*}
f\left(i_{1}\right)+f\left(i_{2}^{\prime}\right) \equiv f\left(i_{1}^{\prime}\right)+f\left(i_{2}\right)(\bmod M) . \tag{7}
\end{equation*}
$$

Since $e_{1}$ is an edge of $G, i_{1} \neq i_{1}^{\prime}$ must hold, and similarly, since $e_{2}$ is an edge of $G, i_{2} \neq i_{2}^{\prime}$ must hold. Moreover, since $e_{1} \neq e_{2}$, either $i_{1} \neq i_{2}$, or $i_{1}^{\prime} \neq i_{2}^{\prime}$ must hold. Combining this with Equation (7), we conclude that both $i_{1} \neq i_{2}$ and $i_{1}^{\prime} \neq i_{2}^{\prime}$ must hold.

We now consider four cases. The first case is when all indices in $\left\{i_{1}, i_{1}^{\prime}, i_{2}, i_{2}^{\prime}\right\}$ are distinct. In this case, we consider the ensemble $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$, where $I=\left\{i_{1}, i_{1}^{\prime}, i_{2}, i_{2}^{\prime}\right\}$, with $x_{i_{1}}=x_{i_{2}^{\prime}}=1$ and $x_{i_{2}}=x_{i_{1}^{\prime}}=-1$. From Equation 7 , we get that $\sum_{i \in I} x_{i} \cdot f(i) \equiv 0(\bmod M)$, so ensemble $\mathcal{B}$ is bad, contradicting the fact that Event $\mathcal{E}^{b}$ did not happen.

The second case is when $i_{1}=i_{2}^{\prime}$ but $i_{1}^{\prime} \neq i_{2}$. Then we construct an ensemble $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$, where $I=\left\{i_{1}, i_{1}^{\prime}, i_{2}\right\}$, with $x_{i_{1}}=2$ and $x_{i_{2}}=x_{i_{1}^{\prime}}=-1$. As before, from Equation 7 , we get that $\sum_{i \in I} x_{i} \cdot f(i) \equiv 0(\bmod M)$, so ensemble $\mathcal{B}$ is bad, contradicting the fact that Event $\mathcal{E}^{b}$ did not happen.

The third case is when $i_{1}^{\prime}=i_{2}$ but $i_{1} \neq i_{2}^{\prime}$. We consider the ensemble $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$, where $I=\left\{i_{1}, i_{1}^{\prime}, i_{2}^{\prime}\right\}$, with $x_{i_{1}}=x_{i_{2}^{\prime}}=1$ and $x_{i_{2}}=-2$. From Equation 7, we get that $\sum_{i \in I} x_{i} \cdot f(i) \equiv$ $0(\bmod M)$, so ensemble $\mathcal{B}$ is bad, contradicting the fact that Event $\mathcal{E}^{b}$ did not happen.

From the above discussion, the only remaining case is when both $i_{1}^{\prime}=i_{2}$ and $i_{1}=i_{2}^{\prime}$ hold. But in this case, $e_{1}=e_{2}$, contradicting our assumption that these two edges are distinct.

Assume that bad event $\mathcal{E}^{b}$ did not happen. For every edge $e \in E(H)$, we denote by $R(e)$ the unique edge of $G$ that serves as the origin edge of $e$. For a subgraph $H^{\prime} \subseteq H$, we let $R\left(H^{\prime}\right)$ be the subgraph of $G$ induced by the set $\left\{R(e) \mid e \in E\left(H^{\prime}\right)\right\}$ of edges; we refer to $R\left(H^{\prime}\right)$ as the origin graph of $H^{\prime}$. In other words, the set of edges of graph $R\left(H^{\prime}\right)$ is $\left\{R(e) \mid e \in E\left(H^{\prime}\right)\right\}$, and the set of its vertices contains every vertex of $G$ that serves as an endpoint to any of these edges.

In the next observation we show that, if bad Event $\mathcal{E}^{b}$ did not happen, then for every cycle $C \subseteq H$ containing at most $q$ edges, every vertex in the corresponding origin-graph $R(C)$ has an even degree.

Observation 5.7. Assume that Event $\mathcal{E}^{b}$ did not happen. Let $C \subseteq H$ be any simple cycle containing at most $q$ edges, and let $G^{\prime}=R(C)$ be the origin graph of $C$. Then every vertex of $G^{\prime}$ has an even degree in $G^{\prime}$.

Proof: Throughout the proof, we assume that Event $\mathcal{E}^{b}$ did not happen, and we fix a simple cycle $C=\left(u_{j_{1}}, \ldots, u_{j_{z}}\right)$ in $H$, with $z \leq q$. For all $1 \leq i \leq z$, we denote $e_{i}=\left(u_{j_{i}}, u_{j_{i+1}}\right)$, and we denote the origin-edge of $e_{i}$ by $e_{i}^{\prime}=R\left(e_{i}\right)=\left(v_{a_{i}}, v_{b_{i}}\right)$. From the definition of graph $H$, there must be an integer $t_{i} \in[M]$, with $j_{i} \equiv f\left(a_{i}\right)+t_{i}(\bmod M)$ and $j_{i+1} \equiv f\left(b_{i}\right)+t_{i}(\bmod M)$. Therefore, for all $1 \leq i \leq z$ :

$$
f\left(a_{i}\right)-f\left(b_{i}\right) \equiv j_{i}-j_{i+1}(\bmod M) .
$$

Summing up the above equality over all $i=1, \ldots z$, we get that

$$
\begin{equation*}
\sum_{1 \leq i \leq z} f\left(a_{i}\right)-\sum_{1 \leq i \leq z} f\left(b_{i}\right) \equiv 0(\bmod M) . \tag{8}
\end{equation*}
$$

Let $A$ be the set of indices lying in $\left\{a_{1}, \ldots, a_{z}\right\}$ (if an index appears several times in $\left\{a_{1}, \ldots, a_{z}\right\}$, we only include it once in $A$ ). For every index $a^{*} \in A$, let $x_{a^{*}}^{\prime}$ be the number of integers $i \in\{1, \ldots, z\}$ with $a_{i}=a^{*}$. Similarly, we let $B$ be the set of indices lying in $\left\{b_{1}, \ldots, b_{z}\right\}$, and for every index $b^{*} \in B$, we let $x_{b^{*}}^{\prime \prime}$ be the number of integers $i \in\{1, \ldots, z\}$ with $b_{i}=a$. We claim that $A=B$ must hold, and, for every index $a \in A, x_{a}^{\prime}=x_{a}^{\prime \prime}$ must hold. Indeed, assume otherwise. We then construct an ensemble $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$ as follows. Set $I$ includes every index $i \in(A \backslash B) \cup(B \backslash A)$; for each such index $i$, we set $x_{i}=x_{i}^{\prime}$ if $i \in A \backslash B$ and $x_{i}=-x_{i}^{\prime \prime}$ otherwise. Additionally, for every index $i \in A \cap B$ with $x_{i}^{\prime} \neq x_{i}^{\prime \prime}$, we include index $i$ in $I$, with $x_{i}=x_{i}^{\prime}-x_{i}^{\prime \prime}$. From our assumptions, $I \neq \emptyset,|I| \leq 2 q$, and for all $i \in I,-q \leq x_{i} \leq q$, with $x_{i} \neq 0$. Therefore, $\mathcal{B}=\left\{I,\left\{x_{i}\right\}_{i \in I}\right\}$ is a valid ensemble. But then, from Equation 8, $\sum_{i \in I} x_{i} \cdot f(i) \equiv 0(\bmod M)$. In other words, ensemble $\mathcal{B}$ is bad, contradicting the assumption that bad event $\mathcal{E}^{b}$ did not happen.
We conclude that $A=B$ must hold, and, for every index $a \in A, x_{a}^{\prime}=x_{a}^{\prime \prime}$ must hold. Therefore, for every vertex $v \in V(G)$, the number of times that $v$ lies in $\left\{v_{a_{1}}, \ldots, v_{a_{z}}\right\}$ is equal to the number of times that $v$ lies in $\left\{v_{b_{1}}, \ldots, v_{b_{z}}\right\}$. Therefore, the number of edges of $\left\{R\left(e_{i}\right) \mid 1 \leq i \leq z\right\}$ that are incident to $v$ is even.

Lastly, we need the following claim.
Claim 5.8. Let $H^{\prime}$ be any subgraph of $H$ with $\left|V\left(H^{\prime}\right)\right| \leq N$. If Event $\mathcal{E}^{b}$ did not happen, then the origin graph $R\left(H^{\prime}\right)$ of $H^{\prime}$ contains at most $c^{*} \cdot\left|V\left(H^{\prime}\right)\right|$ vertices, where $c^{*}$ is a constant independent of $N$.

Proof: Recall that the girth of an unweighted graph $G^{*}$ is the length of the shortest cycle in $G^{*}$. For an integer $t \geq 1$, we say that a subgraph $G^{\prime}$ of $G^{*}$ is a $t$-spanner of $G^{*}$ if $V\left(G^{\prime}\right)=V\left(G^{*}\right)$, and, for every pair $v, v^{\prime}$ of vertices of $G^{*}$, if we denote by $\operatorname{dist}_{G^{*}}\left(v, v^{\prime}\right)$ the length of the shortest $v-v^{\prime}$ path in $G^{*}$, and we define $\operatorname{dist}_{G^{\prime}}\left(v, v^{\prime}\right)$ similarly for $G^{\prime}$, then $\operatorname{dist}_{G^{\prime}}\left(v, v^{\prime}\right) \leq t \cdot \operatorname{dist}_{G^{*}}\left(v, v^{\prime}\right)$.
Consider now any subgraph $H^{\prime}$ of $H$. We use the following algorithm of $\mathrm{ADD}^{+} 93$, whose goal is to construct a $q$-spanner $H^{\prime \prime}$ of $H^{\prime}$ that contains few edges. The algorithm starts with graph $H^{\prime \prime}$, whose vertex set is $V\left(H^{\prime \prime}\right)=V\left(H^{\prime}\right)$, and edge set is empty. It then processes every edge $e \in E\left(H^{\prime}\right)$ one by one. If graph $H^{\prime \prime} \cup\{e\}$ contains a cycle of length at most $q$, then we continue to the next iteration; otherwise, we add $e$ to $H^{\prime \prime}$, and continue to the next iteration. Consider the final graph $H^{\prime \prime}$ that is obtained at the end of the algorithm, once very edge of $H^{\prime}$ is processed. It is immediate to see that the girth of $H^{\prime \prime}$ is greater than $q$. One can also show that the resulting graph $H^{\prime \prime}$ is a $q$-spanner of $H^{\prime}$, but we do not need to use this fact. We use the following theorem from Bol04].

Theorem 5.9 (Theorem 3.7 from (Bol04). Let $G$ be an n-vertex graph with girth greater than $q$, for any integer $q>1$. Then $\left.|E(G)| \leq n \cdot \mid n^{2 /(q-2)}\right\rceil$.

From the above theorem, $\left|E\left(H^{\prime \prime}\right)\right| \leq\left|V\left(H^{\prime}\right)\right|^{1+O(1 / q)}=O\left(\left|V\left(H^{\prime}\right)\right|\right)$, as $\left|V\left(H^{\prime}\right)\right| \leq N$ and $q=$ $\lceil\log N\rceil$.
We denote by $W^{\prime} \subseteq G$ the origin graph of $H^{\prime}$, and we denote by $W^{\prime \prime} \subseteq G$ the origin graph of $H^{\prime \prime}$. Note that $\left|E\left(W^{\prime \prime}\right)\right| \leq\left|E\left(H^{\prime \prime}\right)\right| \leq O\left(\left|V\left(H^{\prime}\right)\right|\right)$, and so $\left|V\left(W^{\prime \prime}\right)\right| \leq O\left(\left|E\left(W^{\prime \prime}\right)\right|\right) \leq O\left(\left|V\left(H^{\prime}\right)\right|\right)$. We next show the following observation.
Observation 5.10. If Event $\mathcal{E}^{b}$ did not happen, then $V\left(W^{\prime}\right)=V\left(W^{\prime \prime}\right)$.

Notice that the observation implies that $\left|V\left(W^{\prime}\right)\right| \leq O\left(\left|V\left(H^{\prime}\right)\right|\right)$, completing the proof of Claim 5.8. It now remains to prove Observation 5.10.

Proof of Observation 5.10. Consider any edge $e \in E\left(H^{\prime}\right) \backslash E\left(H^{\prime \prime}\right)$. From the construction of graph $H^{\prime \prime}$, there must be a simple cycle $C$ in graph $H^{\prime \prime} \cup\{e\}$, whose length is at most $q$. Consider now the subgraph $W_{C}$ of $G$, induced by the edges of $\left\{R\left(e^{\prime}\right) \mid e^{\prime} \in E(C)\right\}$; in other words, $W_{C}=R(C)$. Since the event $\mathcal{E}^{b}$ did not happen, from Observation 5.7, graph $W_{C}$ is an even-degree graph. Therefore, if we denote by $\hat{e}=R(e)$ the origin-edge of $e$, then graph $W_{C} \backslash\{\hat{e}\}$ contains exactly two odd-degree vertices, that serve as endpoints of edge $\hat{e}$ in $G$. Notice however that $W_{C} \backslash\{\hat{e}\} \subseteq W^{\prime \prime}$. Therefore, for every edge $e \in E\left(H^{\prime}\right) \backslash E\left(H^{\prime \prime}\right)$, the endpoints of the origin edge $R(e)$ lie in $W^{\prime \prime}$. It then follows that $W^{\prime}=W^{\prime \prime}$.

Bad Event $\mathcal{E}$. We say that the bad event $\mathcal{E}$ happens if either bad event $\mathcal{E}^{b}$ happens, or bad event $\mathcal{E}^{g}$ does not happen. By using the Union bound, together with Claim 5.3 and Observation 5.5, we get that $\operatorname{Pr}[\mathcal{E}] \leq 0.1$.

We are now ready to complete the proof of Theorem 5.1.

### 5.2 Completing the Reduction from Densest $k$-Subgraph to Dense $k$-Coloring

Let $\operatorname{DkS}(G, k)$ be an input instance of the Densest $k$-Subgraph problem. Denote $V(G)=\left\{v_{1}, \ldots, v_{N}\right\}$. We start by constructing the auxiliary graph $H$, from instance $\operatorname{DkS}(G, k)$. We add isolated vertices to graph $H$, until $|V(H)|$ becomes an integral multiple of $k$, and we denote $|V(H)|=n$. Clearly, $n \leq N^{O(\log N)}$. We then consider instance $\operatorname{DkC}(H, k)$ of the Dense $k$-Coloring problem, where the parameter $k$ remains unchanged. Note that, if Event $\mathcal{E}$ did not happen, then, from the definition of Events $\mathcal{E}$ and $\mathcal{E}^{g}$, there is a collection $\left\{H_{1}, \ldots, H_{r}\right\}$ of $r=\left\lfloor\frac{M}{k \log k}\right\rfloor$ disjoint subsets of vertices of $H$, such that for all $1 \leq j \leq r,\left|V\left(H_{j}\right)\right| \leq k$ holds, and additionally, $\sum_{1 \leq j \leq r}\left|E\left(H_{j}\right)\right| \geq 0.1 \cdot\left\lfloor\frac{M}{k \log k}\right\rfloor \cdot \operatorname{OPT}_{\mathrm{DkS}}(G, k)$. We can then define a solution $\left(S_{1}, \ldots, S_{n / k}\right)$ to instance $\operatorname{DkC}(H, k)$ of the Dense $k$-Coloring problem, as follows. For all $1 \leq i \leq r$, we initially set $S_{i}=V\left(H_{i}\right)$, and for all $r<i \leq n / k$, we set $S_{i}=\emptyset$. Let $U=V(H) \backslash\left(\bigcup_{i=1}^{r} V\left(H_{i}\right)\right)$. Next, we partition the vertices of $U$ by adding them to sets $S_{1}, \ldots, S_{n / r}$ arbitrarily, to ensure that the cardinality of each set is exactly $k$. From the above discussion, if Event $\mathcal{E}$ did not happen, then:

$$
\sum_{i=1}^{n / k}\left|E_{H}\left(S_{i}\right)\right| \geq \sum_{i=1}^{r}\left|E\left(H_{i}\right)\right| \geq 0.1 \cdot\left\lfloor\frac{M}{k \log k}\right\rfloor \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k) \geq \Omega\left(\frac{n}{k \log k}\right) \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k) .
$$

We conclude that, if Event $\mathcal{E}$ did not happen, then $\operatorname{OPT}_{\mathrm{DkC}}(H, k) \geq \Omega\left(\frac{n}{k \log k}\right) \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)$.
We apply the $\alpha(n)$-approximation algorithm for Dense $k$-Coloring to instance $\operatorname{DkC}(H, k)$, and we denote the resulting solution by $\left(U_{1}, \ldots, U_{n / k}\right)$. Note that:

$$
\sum_{i=1}^{n / k}\left|E_{H}\left(U_{i}\right)\right| \geq \frac{\mathrm{OPT}_{\mathrm{DkC}}(H, k)}{\alpha(n)} \geq \Omega\left(\frac{n}{k \cdot \alpha(n) \cdot \log k}\right) \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)
$$

must hold. We let $U \in\left\{U_{1}, \ldots, U_{n / k}\right\}$ be a subset maximizing $\left|E_{H}\left(U_{i}\right)\right|$, so that

$$
\left|E_{H}(U)\right| \geq \frac{\sum_{i=1}^{n / k}\left|E_{H}\left(U_{i}\right)\right|}{n / k} \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{DkS}}(G, k)}{\alpha(n) \cdot \log k}\right)
$$

Let $H^{\prime}=H[U]$, and let $W=R\left(H^{\prime}\right)$ be the origin graph of $H^{\prime}$. From Claim 5.8, if Event $\mathcal{E}$ did not happen, then $|V(W)| \leq c^{*} \cdot\left|V\left(H^{\prime}\right)\right| \leq c^{*} k$, for some universal constant $c^{*}$.
Lastly, we apply the algorithm from Lemma 2.3 to graph $W$, to obtain a subgraph $W^{\prime}$ of $W$ with $\left|V\left(W^{\prime}\right)\right|=k$ and $\left|E\left(W^{\prime}\right)\right| \geq \Omega(|E(W)|)$. We then return $S=\left|V\left(W^{\prime}\right)\right|$ as the solution to the input instance $\operatorname{DkS}(G, k)$ of the Densest $k$-Subgraph problem. From the above discussion, $\left|E_{G}(S)\right| \geq \Omega(|E(W)|) \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{DkS}}(G, k)}{\alpha(n) \cdot \log k}\right) \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{DkS}}(G, k)}{\alpha\left(N^{O(\log N)}\right) \cdot \log N}\right)$. Therefore, if the event $\mathcal{E}$ does not happen, we obtain an $O\left(\alpha\left(N^{O(\log N)}\right) \cdot \log N\right)$-approximate solution to the input instance of the Densest $k$-Subgraph problem. Recall that the probability of Event $\mathcal{E}$ happening is at most 0.1. Lastly, since $|V(H)| \leq N^{O(\log N)}$, it is easy to verify that the running time of the algorithm is at most $N^{O(\log N)}$.

### 5.3 Completing the Reduction from Densest $k$-Subgraph to ( $r, h$ )-Graph Partitioning

Let $\operatorname{DkS}(G, k)$ be an input instance of the Densest $k$-Subgraph problem with $|V(G)|=N$. Our algorithm requires the knowledge of an estimate $h$ on the value of $\mathrm{OPT}_{\mathrm{DkS}}(G, k)$, with $h / 2 \leq$ $\operatorname{OPT}_{\mathrm{DkS}}(G, k) \leq h$. In order to overcome this difficulty, we run the algorithm for every value of $h$ that is an integral power of 2 between 1 and $|E(G)|$, and output the best resulting solution. Therefore, it is now enough to provide a randomized algorithm that, given an estimate $h$ with $h / 2 \leq \mathrm{OPT}_{\mathrm{DkS}}(G, k) \leq h$, with a constant probability produces a solution to instance $\operatorname{DkS}(G, k)$ of Densest $k$-Subgraph whose value is at least $\Omega\left(\frac{\mathrm{OPT}_{\mathrm{Dks}}(G, k)}{\left(\alpha\left(N^{O}(\log N)\right)\right)^{3} \cdot \log ^{2} N}\right)$, such that the running time of the algorithm is $N^{O(\log N)}$. From now on we assume that we are given an integer $h$ with $h / 2 \leq$ $\operatorname{OPT}_{\mathrm{DkS}}(G, k) \leq h$.
As before, we denote $V(G)=\left\{v_{0}, \ldots, v_{N-1}\right\}$, and we construct the auxiliary graph $H$ from instance $\operatorname{DkS}(G, k)$ of Densest $k$-Subgraph. We denote $|V(H)|=n$, so $n \leq N^{O(\log N)}$ holds. We then consider instance $\operatorname{GP}(H, r, h)$ of the $(r, h)$-Graph Partitioning problem, where $r=\left\lfloor\frac{n}{k \log k}\right\rfloor$.
Note that, if Event $\mathcal{E}$ did not happen, then, from the definition of Events $\mathcal{E}$ and $\mathcal{E}^{g}$, there is a collection $\left\{H_{1}, \ldots, H_{r}\right\}$ of $r=\left\lfloor\frac{n}{k \log k}\right\rfloor$ disjoint subgraphs of $H$, such that for all $1 \leq j \leq r$, $\left|E\left(H_{j}\right)\right| \leq \mathrm{OPT}_{\mathrm{DkS}}(G, k) \leq h$ holds, and $\sum_{1 \leq j \leq r}\left|E\left(H_{j}\right)\right| \geq 0.1 \cdot\left\lfloor\frac{n}{k \log k}\right\rfloor \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)$. Therefore, we obtain a solution $\left(H_{1}, \ldots, H_{r}\right)$ to instance $\operatorname{GP}(H, r, h)$ of (r,h)-Graph Partitioning, whose value is at least 0.1 $\left\lfloor\frac{n}{k \log k}\right\rfloor \cdot \operatorname{OPT}_{\mathrm{DkS}}(G, k)$. We conclude that, if Event $\mathcal{E}$ did not happen, then $\mathrm{OPT}_{\mathrm{GP}}(H, r, h) \geq \Omega\left(\frac{n}{k \log k}\right) \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)$.
We apply the $\alpha(n)$-approximation algorithm to instance $\operatorname{GP}(H, r, h)$ of ( $r, \mathrm{~h})$-Graph Partitioning, obtaining a solution $\left(H_{1}^{\prime}, \ldots, H_{r}^{\prime}\right)$, whose value is at least $\frac{\mathrm{OPT}_{\mathrm{GP}}(H, r, h)}{\alpha(n)}$. Denote $\mathcal{H}=\left\{H_{1}^{\prime}, \ldots, H_{r}^{\prime}\right\}$. We partition set $\mathcal{H}$ into two subsets: set $\mathcal{H}^{\prime}$ containing all graphs $H_{j}^{\prime} \in \mathcal{H}$ with $\left|V\left(H^{\prime}\right)\right| \leq 100$.
$\alpha(n) k \log k$, and set $\mathcal{H}^{\prime \prime}$ containing all remaining graphs.
Assume for now that Event $\mathcal{E}$ did not happen. Then, as observed above:

$$
\sum_{H_{j}^{\prime} \in \mathcal{H}}\left|E\left(H_{j}^{\prime}\right)\right| \geq \frac{\mathrm{OPT}_{\mathrm{GP}}(H, r, h)}{\alpha(n)} \geq \frac{n \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)}{20 k \cdot \alpha(n) \cdot \log k}
$$

Clearly, $\left|\mathcal{H}^{\prime \prime}\right| \leq \frac{n}{100 \cdot \alpha(n) \cdot k \log k}$, and so:

$$
\sum_{H_{j}^{\prime} \in \mathcal{H}^{\prime \prime}}\left|E\left(H_{j}^{\prime}\right)\right| \leq \frac{n \cdot h}{100 \cdot \alpha(n) \cdot k \log k} \leq \frac{n \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)}{100 \cdot \alpha(n) \cdot k \log k} .
$$

Altogether, we get that, if Event $\mathcal{E}$ did not happen, then:

$$
\sum_{H_{j}^{\prime} \in \mathcal{H}^{\prime}}|E(H)| \geq \frac{n \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)}{50 k \cdot \alpha(n) \cdot \log k}
$$

We let $H^{*} \in \mathcal{H}^{\prime}$ be the graph maximizing the number of edges. Since $\left|\mathcal{H}^{\prime}\right| \leq|\mathcal{H}|=r=\left\lfloor\frac{n}{k \log k}\right\rfloor$, from the above discussion, if Event $\mathcal{E}$ did not happen, then:

$$
\left|E\left(H^{*}\right)\right| \geq \frac{n \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)}{r \cdot 50 k \cdot \alpha(n) \cdot \log k} \geq \frac{\mathrm{OPT}_{\mathrm{DkS}}(G, k)}{50 \alpha(n)}
$$

From the definition of the collection $\mathcal{H}^{\prime}$ of graphs, $\left|V\left(H^{*}\right)\right| \leq 100 \cdot \alpha(n) k \log k$.
Let $W=R\left(H^{*}\right)$ be the origin graph of $H^{*}$. From Claim 5.8, if Event $\mathcal{E}$ did not happen, then $|V(W)| \leq c^{*} \cdot\left|V\left(H^{*}\right)\right| \leq O(\alpha(n) k \log k)$.
Lastly, we apply the algorithm from Lemma 2.3 to graph $W$, to obtain a subgraph $W^{\prime}$ of $W$ with $\left|V\left(W^{\prime}\right)\right| \leq k$ and $\left|E\left(W^{\prime}\right)\right| \geq \Omega\left(\frac{|E(W)|}{(\alpha(n))^{2} \cdot \log ^{2} k}\right)$. We then return $S=\left|V\left(W^{\prime}\right)\right|$ as the solution to the input instance $\operatorname{DkS}(G, k)$ of the Densest $k$-Subgraph problem. From the above discussion, $|V(S)| \leq k$, and, if Event $\mathcal{E}$ did not happen, then:

$$
\begin{aligned}
\left|E_{G}(S)\right| & \geq \Omega\left(\frac{|E(W)|}{(\alpha(n))^{2} \cdot \log ^{2} k}\right) \\
& \geq \Omega\left(\frac{\left|E\left(H^{*}\right)\right|}{(\alpha(n))^{2} \cdot \log ^{2} k}\right) \\
& \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{DkS}}(G, k)}{(\alpha(n))^{3} \cdot \log ^{2} k}\right) \\
& \geq \Omega\left(\frac{\mathrm{OPT}_{\mathrm{DkS}}(G, k)}{\left(\alpha\left(N^{O(\log N)}\right)\right)^{3} \cdot \log ^{2} N}\right) .
\end{aligned}
$$

Therefore, if the event $\mathcal{E}$ does not happen, we obtain an $O\left(\left(\alpha\left(N^{O(\log N)}\right)\right)^{3} \cdot \log ^{2} N\right)$-approximate solution to the input instance of the Densest $k$-Subgraph problem. Recall that the probability of

Event $\mathcal{E}$ happening is at most 0.1 . Lastly, since $|V(H)| \leq N^{O(\log N)}$, it is easy to verify that the running time of the algorithm is at most $N^{O(\log N)}$.

## 6 Reductions between (r,h)-Graph Partitioning and Maximum BoundedCrossing Subgraph

In this section we establish a connection between the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning and Maximum BoundedCrossing Subgraph problems, by proving the following two theorems.

Theorem 6.1. Let $\alpha: \mathbb{Z}^{+} \rightarrow \mathbb{Z}^{+}$be an increasing function with $\alpha(n)=o(n)$. Assume that there exists an efficient $\alpha(n)$-approximation algorithm for the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning problem, where $n$ is the number of vertices in the input graph. Then there exists an efficient $O(\alpha(N) \cdot \operatorname{poly} \log N)$ approximation algorithm for Maximum Bounded-Crossing Subgraph, where $N$ is the number of vertices in the input instance of Maximum Bounded-Crossing Subgraph.

Theorem 6.2. Let $\alpha: \mathbb{Z}^{+} \rightarrow \mathbb{Z}^{+}$be an increasing function with $\alpha(n)=o(n)$. Assume that there exists an efficient $\alpha(N)$-approximation algorithm for the Maximum Bounded-Crossing Subgraph problem, where $N$ is the number of vertices in the input graph. Then there exists an efficient $O\left((\alpha(n))^{2} \cdot\right.$ poly $\left.\log n\right)$-approximation algorithm for $(\mathrm{r}, \mathrm{h})$-Graph Partitioning, where $n$ is the number of vertices in the input instance of (r,h)-Graph Partitioning.

By combining Theorem 6.2 with Corollary 5.2, we obtain the following corollary.
Corollary 6.3. Assume that Conjecture 2 holds and that NP $\nsubseteq \operatorname{BPTIME}\left(n^{O(\log n)}\right)$. Then for some constant $0<\varepsilon^{\prime} \leq 1 / 2$, there is no efficient $2^{(\log n)^{\varepsilon^{\prime}}}$-approximation algorithm for Maximum Bounded-Crossing Subgraph.

Proof: Assume that Conjecture 2 holds and that NP $\nsubseteq \operatorname{DTIME}\left(n^{O(\log n)}\right)$. Then, from Corollary 5.2, for some constant $0<\varepsilon \leq 1 / 2$, there is no efficient factor- $2^{(\log n)^{\varepsilon}}$-approximation algorithm for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, where $n$ is the number of vertices in the input graph.

We let $\varepsilon^{\prime}=\varepsilon / c$, where $c$ is a sufficiently large constant. We now prove that there is no efficient $2^{(\log n)^{\varepsilon^{\prime}}}$-approximation algorithm for Maximum Bounded-Crossing Subgraph. Indeed, assume for contradiction that there is an efficient $2^{(\log n)^{\varepsilon^{\prime}}}$-approximation algorithm $\mathcal{A}$ for Maximum BoundedCrossing Subgraph. From Theorem 6.2, there is an efficient $c^{\prime} \cdot(\alpha(N))^{2} \cdot$ poly $\log N$-approximation algorithm for ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, where $N$ is the number of vertices in the input instance of $(r, h)$-Graph Partitioning, $c^{\prime}$ is some constant, and $\alpha(N)=2^{(\log N)^{\varepsilon^{\prime}}}$. Notice however that $c^{\prime} \cdot(\alpha(N))^{2}$. poly $\log N \leq 2^{(\log N)^{\varepsilon}}$ holds, if the constant $c$ is large enough.
Therefore, we obtain an efficient factor- $2^{(\log n)^{\varepsilon}}$-approximation algorithm for (r,h)-Graph Partitioning, a contradiction.
In the remainder of this section, we prove Theorems 6.2 and 5.2. We start by proving two auxiliary lemmas that will be used in the proofs of both theorems. We then complete the proofs of Theorem 6.1 and Theorem 6.2 in sections Section 6.2 and Section 6.3 , respectively.

### 6.1 Auxiliary Lemmas

We start with the following definition, that will be used in the proofs of both theorems.
Definition 6.4. Let $\operatorname{GP}(G, r, h)$ be an instance of (r,h)-Graph Partitioning, and let $\left\{H_{1}, \ldots, H_{r}\right\}$ be a solution to this instance. We say that this solution is good, if for all $1 \leq i \leq r, h / 2 \leq\left|E\left(H_{i}\right)\right| \leq h$.

We are now ready to state the first auxiliary lemma.
Lemma 6.5. There is an efficient algorithm, that, given a graph $G$ with $|V(G)|=n$, integers $r, h>0$ and any solution $\mathcal{H}$ to instance $\operatorname{GP}(G, r, h)$ of $(r, h)$-Graph Partitioning, computes positive integers $r^{*} \leq r, h^{*} \leq h$ and a subset $\mathcal{H}^{*} \subseteq \mathcal{H}$ of subgraphs of $G$, such that $\mathcal{H}^{*}$ is a good solution to instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$, of $(r, h)$-Graph Partitioning, and $\sum_{H \in \mathcal{H}^{*}}|E(H)| \geq \frac{\sum_{H^{\prime} \in \mathcal{H}}\left|E\left(H^{\prime}\right)\right|}{4 \log n}$.

Proof: The idea of the proof is to partition the graphs in $\mathcal{H}$ geometrically into groups by the cardinalities of their edge sets, and then select a group maximizing the total number of edges in its subgraphs.
Specifically, let $q=\lceil 2 \log n\rceil$. For all $1 \leq i \leq q$, we let $\mathcal{H}_{i} \subseteq \mathcal{H}$ contain all graphs $H$ with $2^{i-1} \leq|E(H)|<2^{i}$. It is easy to verify that $\mathcal{H}_{1} \ldots, \mathcal{H}_{q}$ partition $\mathcal{H}$. Clearly, there must be an index $1 \leq i^{*} \leq q$, with $\sum_{H \in \mathcal{H}_{i^{*}}}|E(H)| \geq \frac{\sum_{H^{\prime} \in \mathcal{H}}\left|E\left(H^{\prime}\right)\right|}{q} \geq \frac{\sum_{H^{\prime} \in \mathcal{H}}\left|E\left(H^{\prime}\right)\right| \text {. We set } r^{*}=\left|\mathcal{H}_{i^{*}}\right|, h^{*}=2^{i^{*}} \text {, and } \text {, } \log n}{}$. we let $\mathcal{H}^{*}=\mathcal{H}_{i^{*}}$. It is immediate to verify that $\mathcal{H}^{*}$ is a good solution to instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$, of $(r, h)$-Graph Partitioning, and, from the above discussion, $\sum_{H \in \mathcal{H}^{*}}|E(H)| \geq \frac{\sum_{H^{\prime} \in \mathcal{H}}\left|E\left(H^{\prime}\right)\right|}{4 \log n}$.
We are now ready to prove our second auxiliary lemma.
Lemma 6.6. There is an efficient algorithm, whose input consists of an instance $\operatorname{MBCS}(G, L)$ of the Maximum Bounded-Crossing Subgraph problem with $|V(G)|=N$, where $N$ is greater than a sufficiently large constant, together with a solution $H$ to this instance, such that $|E(H)| \geq$ $4 N \log ^{6} N$ holds. The algorithm computes integers $r, h>0$, such that $r \cdot h^{2} \leq L \cdot \log ^{6} N$ and $r \cdot h \geq \Omega(|E(H)| / \log N)$ hold, together with a good solution $\left\{H_{1}, \ldots, H_{r}\right\}$ to instance $\operatorname{GP}(G, r, h)$ of ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning.

Proof: Let $\hat{G}$ be any graph. A cut in $\hat{G}$ is a partition $(A, B)$ of vertices of $\hat{G}$ into two non-empty subsets. The value of the cut $(A, B)$ is $|E(A, B)|$. For a parameter $1 / 2<\beta<1$, we say that cut $(A, B)$ is $\beta$-balanced, if $|E(A)|,|E(B)| \leq \beta \cdot|E(\hat{G})|$. We say that cut $(A, B)$ is a minimum $\beta$-balanced cut if it is a $\beta$-balanced cut whose value is the smallest among all such cuts. We use the following theorem that follows from the results of ARV09], and was formally proved in [CT22].

Theorem 6.7 (Theorem 4.11 in the full version of [CT22]). There is an efficient algorithm, that, given a graph $\hat{G}$ with $|V(\hat{G})|=\hat{N}$, computes a $\gamma$-balanced cut in $\hat{G}$, whose value is at most $O(\sqrt{\log \hat{N}})$ times the value of the minimum (3/4)-balanced cut in $\hat{G}$, for some universal constant $3 / 4<\gamma<1$ that does not depend on $\hat{N}$.

We use the following theorem, that is a simple corollary of the Planar Separator Theorem by Lipton and Tarjan LT79, and was formally proved in CT22. A variation of this theorem for vertex-balanced cuts was proved in PSS96.

Theorem 6.8 (Lemma 4.12 in the full version of $[\mathrm{CT} 22]$ ). Let $\hat{G}$ be a connected graph with $m$ edges and maximum vertex degree $\Delta<\frac{m}{2^{40}}$. If $\operatorname{CrN}(\hat{G}) \leq \frac{m^{2}}{2^{40}}$, then the value of the minimum (3/4)-balanced cut in $\hat{G}$ is at most $O(\sqrt{\operatorname{CrN}(\hat{G})+\Delta \cdot m})$.

We are now ready to complete the proof of Lemma 6.6. Recall that we are given an instance $\operatorname{MBCS}(G, L)$ of Maximum Bounded-Crossing Subgraph, where $|V(G)|=N$, together with a solution $H$ to this instance, such that $|E(H)| \geq 4 N \log ^{6} N$.
The algorithm starts by iteratively decomposing graph $H$ into smaller subgraphs. Throughout the decomposition procedure, we maintain a collection $\mathcal{H}$ of connected subgraphs of $H$, that are all mutually disjoint. Each graph $H^{\prime} \in \mathcal{H}$ is marked as either active or inactive. At the beginning of the algorithm, we let $\mathcal{H}$ contain all connected components of $H$, which are all marked as active. The algorithm performs iterations, as long as at least one graph in $\mathcal{H}$ is inactive.
In order to execute an iteration, we select an arbitrary active graph $H^{\prime} \in \mathcal{H}$. We apply the algorithm from Theorem 6.7 to compute a $\gamma$-balanced cut $(A, B)$ of $H^{\prime}$. If $\left|E_{H^{\prime}}(A, B)\right| \geq \frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2} N}$, then we mark $H^{\prime}$ as inactive and continue to the next iteration. Otherwise, we remove graph $H^{\prime}$ from $\mathcal{H}$, and we add all connected components of graphs $H^{\prime}[A]$ and $H^{\prime}[B]$ to $\mathcal{H}$, that are all marked as active graphs. We then continue to the next iteration. This completes the description of the decomposition procedure. Let $\mathcal{H}^{\prime}$ be the collection $\mathcal{H}$ of subgraphs of $H$ that we obtain at the end of the procedure. We prove the following simple observation.
Observation 6.9. $\sum_{H^{\prime} \in \mathcal{H}^{\prime}}\left|E\left(H^{\prime}\right)\right| \geq|E(H)| / 2$.
Proof: We use a charging scheme. We observe the set $\mathcal{H}$ of graphs over the course of the partitioning procedure. Throughout the execution of the partitioning procedure, we denote by $E^{\prime}=E(H) \backslash$ $\left(\bigcup_{H^{\prime} \in \mathcal{H}} E\left(H^{\prime}\right)\right)$, and we call the edges of $E^{\prime}$ deleted edges. Over the course of the partitioning procedure we maintain, for every edge $e \in E(H)$, a non-negative value $c(e)$, that we refer to as the charge of $e$. We will ensure that, at every point of the algorithm's execution, $\sum_{e \in E(H)} c(e) \geq\left|E^{\prime}\right|$, and that, for every edge $e \in E(H), c(e) \leq 1 / 2$ always holds. We note that, even when an edge $e$ is added to the set $E^{\prime}$ of deleted edges, its charge $c(e)$ may remain strictly positive. It is then easy to verify that, at the end of the algorithm, $\left|E^{\prime}\right| \leq \sum_{e \in E(H)} c(e) \leq|E(H)| / 2$ holds, and so $\sum_{H^{\prime} \in \mathcal{H}^{\prime}}\left|E\left(H^{\prime}\right)\right| \geq|E(H)| / 2$.
It now remains to describe the assignment of the charge values $c(e)$ to the edges $e \in E(H)$, for which the above properties hold. Initially, $E^{\prime}=\emptyset$, and we set $c(e)=0$ for every edge $e \in E(H)$.
Consider now some iteration of the algorithm, and assume that, at the beginning of the iteration, $\sum_{e \in E(H)} c(e) \geq\left|E^{\prime}\right|$ holds. Let $H^{\prime} \in \mathcal{H}$ be the graph that was processed in the current iteration, and let $(A, B)$ be the cut in $H^{\prime}$ that the algorithm computed. If $\left|E_{H^{\prime}}(A, B)\right| \geq \frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2} N}$, then no new edges were added to $E^{\prime}$ in the current iteration, and the charge values $c(e)$ remain unchanged for all edges $e \in E(H)$. Assume now that $\left|E_{H^{\prime}}(A, B)\right|<\frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2} N}$ holds, and denote $E^{\prime \prime}=E_{H^{\prime}}(A, B)$. Then in the current iteration, the edges of $E^{\prime \prime}$ were added to set $E^{\prime}$. We increase the charge $c(e)$ of every edge $e \in E\left(H^{\prime}\right)$ by $\frac{\left|E^{\prime \prime}\right|}{\left|E\left(H^{\prime}\right)\right|}$, and leave all other edge charges unchanged. This ensures that $\sum_{e \in E(H)} c(e) \geq\left|E^{\prime}\right|$ holds at the end of the iteration. Since $\left|E^{\prime \prime}\right|<\frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2} N}$, for every edge $e \in E\left(H^{\prime}\right)$, the charge $c(e)$ increases by at most $\frac{1}{\log ^{2} N}$ in the current iteration.

From the above discussion, at the end of the algorithm, $\left|E^{\prime}\right| \leq \sum_{e \in E(H)} c(e)$ holds. It now remains to show that for every edge $e \in E(H), c(e) \leq 1 / 2$ holds at the end of the algorithm.
Consider any edge $e \in E(H)$, and denote by $H_{1}, H_{2}, \ldots, H_{r}$ the sequence of subgraphs of $H$ that belonged to $\mathcal{H}$ over the course of the algorithm, and contained $e$. In other words, $H_{1}=H$, and, for all $1<i \leq r$, graph $H_{i}$ was obtained via a balanced cut from graph $H_{i-1}$. Then the charge of $e$ has increased in at most $r+1$ iterations, and in each such iteration, the increase in the charge was bounded by $\frac{1}{\log ^{2} N}$. Furthermore, for all $1<i \leq r,\left|E\left(H_{i}\right)\right| \leq \gamma \cdot\left|E\left(H_{i+1}\right)\right|$ holds, and so $r \leq O(\log N)$ as $\gamma$ is a constant. Therefore, at the end of the algorithm, $c(e) \leq \frac{r+1}{\log ^{2} N} \leq \frac{1}{2}$, since we have assumed that $N$ is sufficiently large.

Consider now the final collection $\mathcal{H}^{\prime}$ of graphs. We say that a graph $H^{\prime} \in \mathcal{H}^{\prime}$ is dense iff $\left|E\left(H^{\prime}\right)\right| \geq$ $\left|V\left(H^{\prime}\right)\right| \cdot \log ^{6} N$, and otherwise we say it is sparse. We partition the set $\mathcal{H}^{\prime}$ of graphs into a collection $\mathcal{H}^{d}$ containing all dense graphs and a collection $\mathcal{H}^{s}$ containing all sparse graphs. We need the following simple observation.
Observation 6.10. $\sum_{H^{\prime} \in \mathcal{H}^{d}}\left|E\left(H^{\prime}\right)\right| \geq \frac{|E(H)|}{4} \geq N \log ^{6} N$.
Proof: Assume for contradiction that $\sum_{H^{\prime} \in \mathcal{H}^{d}}\left|E\left(H^{\prime}\right)\right|<\frac{|E(H)|}{4}$. Since, from Observation 6.9 $\sum_{H^{\prime} \in \mathcal{H}^{\prime}}\left|E\left(H^{\prime}\right)\right| \geq|E(H)| / 2$, we get that $\sum_{H^{\prime} \in \mathcal{H}^{s}}\left|E\left(H^{\prime}\right)\right|>\frac{|E(H)|}{4}$. However:

$$
\sum_{H^{\prime} \in \mathcal{H}^{s}}\left|E\left(H^{\prime}\right)\right| \leq \sum_{H^{\prime} \in \mathcal{H}^{s}}\left|V\left(H^{\prime}\right)\right| \cdot \log ^{6} N \leq N \log ^{6} N
$$

We then conclude that $|E(H)|<4 \sum_{H^{\prime} \in \mathcal{H}^{s}}\left|E\left(H^{\prime}\right)\right| \leq 4 N \log ^{6} N$, contradicting the statement of Lemma 6.6.

We also need the following obsevation.
Observation 6.11. Let $H^{\prime} \in \mathcal{H}^{d}$ be a dense graph. Then $\operatorname{CrN}\left(H^{\prime}\right) \geq \Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|^{2}}{\log ^{5} N}\right)$.
We provide the proof of Observation 6.11 below, after we complete the proof of Lemma 6.6 using it. Let $r^{\prime}=\left|\mathcal{H}^{d}\right|$ and $h^{\prime}=\max _{H^{\prime} \in \mathcal{H}^{d}}\left\{\left|E\left(H^{\prime}\right)\right|\right\}$. Clearly, collection $\mathcal{H}^{d}$ of graphs is a valid solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of $(r, h)$-Graph Partitioning. We apply the algorithm from Lemma 6.5 to the soluton $\mathcal{H}^{d}$ to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of $(r, h)$-Graph Partitioning. Recall that the algorithm computes positive integers $r \leq r^{\prime}, h \leq h^{\prime}$, and a subset $\mathcal{H}^{*} \subseteq \mathcal{H}^{d}$ of subgraphs of $G$, such that $\mathcal{H}^{*}$ is a good solution to instance $\operatorname{GP}(G, r, h)$, of $(r, h)$-Graph Partitioning, and $\sum_{H^{\prime} \in \mathcal{H}^{*}}\left|E\left(H^{\prime}\right)\right| \geq \frac{\sum_{H^{\prime} \in \mathcal{H}^{d}}\left|E\left(H^{\prime}\right)\right|}{4 \log N}$. It now remains to verify that $r \cdot h^{2} \leq L \cdot \log ^{6} N$ and $r \cdot h \geq \Omega(|E(H)| / \log N)$ hold.

Observe first that:

$$
r \cdot h \geq \sum_{H^{\prime} \in \mathcal{H}^{*}}\left|E\left(H^{\prime}\right)\right| \geq \frac{\sum_{H^{\prime} \in \mathcal{H}^{d}}\left|E\left(H^{\prime}\right)\right|}{4 \log N} \geq \frac{|E(H)|}{16 \log N}
$$

from Observation 6.10.
Finally, since for each graph $H^{\prime} \in \mathcal{H}^{d}, \operatorname{CrN}\left(H^{\prime}\right) \geq \Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|^{2}}{\log ^{5} N}\right)$ holds, and since, for every graph $H^{\prime} \in \mathcal{H}^{*},\left|E\left(H^{\prime}\right)\right| \geq h / 2$ holds, we get that:

$$
r \cdot h^{2} \leq \sum_{H^{\prime} \in \mathcal{H}^{*}} 4 \cdot\left|E\left(H^{\prime}\right)\right|^{2} \leq \sum_{H^{\prime} \in \mathcal{H}^{*}} O\left(\log ^{5} N\right) \cdot \mathrm{CrN}\left(H^{\prime}\right) \leq O\left(\log ^{5} N\right) \cdot \mathrm{CrN}(H) \leq L \cdot \log ^{6} N
$$

since $N$ is large enough.
In order to complete the proof of Lemma 6.6, it is now enough to prove Observation 6.11, which we do next.
Proof of Observation 6.11. Since graph $H^{\prime}$ is marked inactive by the algorithm, the $\gamma$-balanced cut of $H^{\prime}$ computed by the algorithm from Theorem 6.7 had value at least $\frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2} N}$. Therefore the minimum (3/4)-balanced cut of $H^{\prime}$ has value at least $\Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2.5} N}\right)$.
Let $\Delta$ denote the maximum vertex degree in $H^{\prime}$. Clearly, $\Delta \leq V\left(H^{\prime}\right)$ must hold. On the other hand, from the definition of a dense graph, $\left|E\left(H^{\prime}\right)\right| \geq\left|V\left(H^{\prime}\right)\right| \cdot \log ^{6} N$, and so $\Delta \leq\left|V\left(H^{\prime}\right)\right| \leq$ $\frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{G} N}<\frac{\left|E\left(H^{\prime}\right)\right|}{2^{40}}$, if $N$ is sufficiently large.
Recall that, from Theorem 6.8, either $\operatorname{CrN}\left(H^{\prime}\right) \geq \frac{\left|E\left(H^{\prime}\right)\right|^{2}}{2^{40}}$, or the value of minimum (3/4)-balanced cut in $H^{\prime}$ is at most $\sqrt{\operatorname{CrN}\left(H^{\prime}\right)+\Delta \cdot\left|E\left(H^{\prime}\right)\right| \text {. In the former case, we immediately get that }}$ $\mathrm{CrN}\left(H^{\prime}\right) \geq \Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|^{2}}{\log ^{5} N}\right)$. In the latter case, since the value of the minimum (3/4)-balanced cut in $H^{\prime}$ is $\Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2.5} N}\right)$, we get that $\sqrt{\operatorname{CrN}\left(H^{\prime}\right)+\Delta \cdot\left|E\left(H^{\prime}\right)\right|} \geq \Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{2.5} N}\right)$. Moreover, since $\Delta \leq \frac{\left|E\left(H^{\prime}\right)\right|}{\log ^{6} N}$ :

$$
\operatorname{CrN}\left(H^{\prime}\right) \geq \Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|^{2}}{\log ^{5} N}\right)-\Delta \cdot\left|E\left(H^{\prime}\right)\right| \geq \Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|^{2}}{\log ^{5} N}\right)-\frac{\left|E\left(H^{\prime}\right)\right|^{2}}{\log ^{6} N} \geq \Omega\left(\frac{\left|E\left(H^{\prime}\right)\right|^{2}}{\log ^{5} N}\right)
$$

since $N$ is sufficiently large.

### 6.2 Reduction from Maximum Bounded-Crossing Subgraph to ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning: Proof of Theorem 6.1

In this subsection we prove Theorem 6.1. Let $\operatorname{MBCS}(G, L)$ be a given instance of Maximum Bounded-Crossing Subgraph, with $|V(G)|=N$. Note that we can assume without loss of generality that $G$ contains no isolated vertices, since all such vertices can be deleted without changing the problem.
As our first step, we compute an arbitrary spanning forest $F$ of graph $G$. Since $G$ contains no isolated vertices, $|E(F)| \geq N / 2$. Clearly, $\operatorname{CrN}(F)=0$. Consider now the optimal solution $H^{*}$ to instance $\operatorname{MBCS}(G, L)$ of Maximum Bounded-Crossing Subgraph. If $\left|E\left(H^{*}\right)\right|<4 N \cdot \log ^{6} N$, then $F$ is a factor- $O\left(\log ^{6} N\right)$ approximate solution to instance $\operatorname{MBCS}(G, L)$.
Assume now that $\left|E\left(H^{*}\right)\right| \geq 4 N \log ^{6} N$. Then, from Lemma 6.6, there exist integers $r, h>0$ with $r \cdot h^{2} \leq L \cdot \log ^{6} N$ and $r \cdot h \geq \Omega\left(\left|E\left(H^{*}\right)\right| / \log N\right)$, such that there exists a good solution to instance $\operatorname{GP}(G, r, h)$ of (r,h)-Graph Partitioning. We will now attempt to guess such integers $r, h$, and then use the approximation algorithm for the ( $r, h$ )-Graph Partitioning problem, in order to compute a solution to the corresponding instance $\operatorname{GP}(G, r, h)$ of $(r, h)$-Graph Partitioning, whose value is sufficiently high.

We say that a pair $(r, h)$ of positive integers is eligible, if $r \cdot h^{2} \leq L \cdot \log ^{6} N$. For each eligible pair $(r, h)$ of integers, we consider the instance $\operatorname{GP}(G, r, h)$ of (r,h)-Graph Partitioning and we use the $\alpha(n)$-approximation algorithm for the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning problem to compute a solution $\mathcal{H}_{r, h}$ to the instance $\mathrm{GP}(G, r, h)$, such that $\sum_{H^{\prime} \in \mathcal{H}_{r, h}}\left|E\left(H^{\prime}\right)\right| \geq \frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{\alpha(N)}$ (since $\left.|V(G)|=N\right)$. We use the following observation.

Observation 6.12. If $\operatorname{OPT}_{\mathrm{MBCS}}(G, L) \geq 4 N \log ^{6} N$, then there exists an eligible pair $(r, h)$ of integers, with:

$$
\sum_{H^{\prime} \in \mathcal{H}_{r, h}}\left|E\left(H^{\prime}\right)\right| \geq \Omega\left(\frac{\operatorname{OPT}_{\mathrm{MBCS}}(G, L)}{\alpha(N) \log N}\right)
$$

Proof: Let $H^{*}$ be an optimal solution to the instance $\operatorname{MBCS}(G, L)$. From Lemma 6.6, there is an eligible pair $\left(r^{*}, h^{*}\right)$ of integers, with $r^{*} \cdot h^{*} \geq \Omega\left(E\left(H^{*}\right) \mid / \log N\right)$, so that there exists a good solution $\mathcal{H}^{*}$ to instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$ of (r,h)-Graph Partitioning. From the definition of a good solution,

$$
\sum_{H^{\prime} \in \mathcal{H}^{*}}\left|E\left(H^{\prime}\right)\right| \geq \frac{r^{*} h^{*}}{2} \geq \Omega\left(\frac{\left|E\left(H^{*}\right)\right|}{\log N}\right) \geq \Omega\left(\frac{\operatorname{OPT}_{\mathrm{MBCS}}(G, L)}{\log N}\right)
$$

Therefore, if $\mathcal{H}_{r^{*}, h^{*}}$ is the approximate solution that we obtained for instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$ of (r,h)-Graph Partitioning, then:

$$
\sum_{H^{\prime} \in \mathcal{H}_{r^{*}, h^{*}}}\left|E\left(H^{\prime}\right)\right| \geq \Omega\left(\frac{\operatorname{OPT}_{\mathrm{GP}}\left(G, r^{*}, h^{*}\right)}{\alpha(N)}\right) \geq \Omega\left(\frac{\sum_{H^{\prime} \in \mathcal{H}^{*}}\left|E\left(H^{\prime}\right)\right|}{\alpha(N)}\right) \geq \Omega\left(\frac{\operatorname{OPT}_{\mathrm{MBCS}}(G, L)}{\alpha(N) \log N}\right)
$$

and the observation follows.
Let $\left(r^{\prime}, h^{\prime}\right)$ be the eligible pair of integers that maximizes $\sum_{H \in \mathcal{H}_{r^{\prime}, h^{\prime}}}|E(H)|$. For every graph $H \in \mathcal{H}_{r^{\prime}, h^{\prime}}$, let $\tilde{H}$ be a graph that is obtained from $H$ as follows. We set $V(\tilde{H})=V(H)$, and we let $E(\tilde{H})$ contain an arbitrary subset of $\left\lfloor\frac{|E(H)|}{\log ^{3} N}\right\rfloor$ edges of $E(H)$. Note that:

$$
\frac{|E(H)|}{2 \log ^{3} N} \leq|E(\tilde{H})| \leq \frac{h^{\prime}}{\log ^{3} N}
$$

Finally, we define a graph $H^{\prime}=\bigcup_{H \in \mathcal{H}_{r^{\prime}, h^{\prime}}} \tilde{H}$. Since the crossing number of any $m$-edge graph is bounded by $m^{2}$, it is easy to verify that:

$$
\operatorname{CrN}\left(H^{\prime}\right) \leq \sum_{H \in \mathcal{H}_{r^{\prime}, h^{\prime}}} \operatorname{CrN}(\tilde{H}) \leq r^{\prime} \cdot \frac{\left(h^{\prime}\right)^{2}}{\log ^{6} N} \leq L
$$

Moreover:

$$
\left|E\left(H^{\prime}\right)\right|=\sum_{H \in \mathcal{H}_{r^{\prime}, h^{\prime}}}|E(\tilde{H})| \geq \sum_{H \in \mathcal{H}_{r, h}} \frac{|E(H)|}{\log ^{3} N} \geq \Omega\left(\frac{\operatorname{OPT}_{\mathrm{MBCS}}(G, L)}{\alpha(N) \log ^{4} N}\right)
$$

Recall that we have computed a spanning forest $F$ of $G$. We return the graph in $\left\{F, H^{\prime}\right\}$ that contains more edges as the outcome of the algorithm. From the above discussion, we obtain an $O(\alpha(N) \cdot$ poly $\log N)$-approximate solution.

### 6.3 Reduction from (r,h)-Graph Partitioning to Maximum Bounded-Crossing SubgraphProof of Theorem 6.2

In this subsection we prove Theorem 6.2. Let $\operatorname{GP}(G, r, h)$ be the input instance of ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, and denote $n=|V(G)|$. For convenience, we will assume that the value $C^{*}=$ $\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$ is known to the algorithm: since $0 \leq \operatorname{OPT}_{\mathrm{GP}}(G, r, h) \leq|E(G)|$ holds, and $\mathrm{OPT}_{\mathrm{GP}}(G, r, h)$ is an integer, we can try all possible guesses for the value $C^{*}$, and then output the best of the resulting solutions. It is sufficient to ensure that the algorithm correctly computes an $O\left((\alpha(n))^{2} \cdot\right.$ poly $\left.\log n\right)$-approximate solution to instance $\mathrm{GP}(G, r, h)$ if the guess $C^{*}$ is correct, that is, $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$. From now on we assume that we are given a value $C^{*}=\mathrm{OPT}_{\mathrm{GP}}(G, r, h)$.
We distinguish between two cases. The first case happens if $C^{*} \geq 16 n \alpha(n) \log ^{7} n$. In this case, we proceed as follows. By applying Lemma 6.5 to the optimal solution to instance $\operatorname{GP}(G, r, h)$, we conclude that there are positive integers $r^{*} \leq r, h^{*} \leq h$ and a collectoin $\mathcal{H}^{*}$ of subgraphs of $G$, such that $\mathcal{H}^{*}$ is a good solution to instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$, of (r,h)-Graph Partitioning, and $r^{*} h^{*} \geq \sum_{H \in \mathcal{H}^{*}}|E(H)| \geq \frac{C^{*}}{4 \log n}$. Since the values of integers $r^{*}, h^{*}$ are not known to our algorithm, we will try all possible candidate values $1 \leq r^{\prime} \leq r$ and $1 \leq h^{\prime} \leq h$ with $r^{\prime} h^{\prime} \geq \frac{C^{*}}{4 \log n}$. For each such pair ( $r^{\prime}, h^{\prime}$ ) of integers, we will compute a solution $\mathcal{H}_{r^{\prime}, h^{\prime}}$ to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$. We will then output the best solution from among $\left\{\mathcal{H}_{r^{\prime}, h^{\prime}} \left\lvert\, r^{\prime} h^{\prime} \geq \frac{C^{*}}{4 \log n}\right.\right\}$. It is sufficient to ensure that, for integers $\left(r^{\prime}, h^{\prime}\right)=\left(r^{*}, h^{*}\right)$, the value of the resulting solution $\mathcal{H}_{r^{\prime}, h^{\prime}}$ is close to $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$.
Consider now a pair of integers $1 \leq r^{\prime} \leq r$ and $1 \leq h^{\prime} \leq h$ with $r^{\prime} h^{\prime} \geq \frac{C^{*}}{4 \log n}$, and assume that values $r^{\prime}, h^{\prime}$ were guessed correctly, that is, $\left(r^{\prime}, h^{\prime}\right)=\left(r^{*}, h^{*}\right)$. In other words, there is a good solution $\mathcal{H}^{*}$ to instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$, of (r,h)-Graph Partitioning, whose value is at least $\frac{C^{*}}{4 \log n}$. Consider now the graph $H^{\prime}=\bigcup_{H \in \mathcal{H}^{*}} H$. Since the crossing number of a graph $H$ may not be higher than $|E(H)|^{2}$, we get that $\operatorname{CrN}(H) \leq r^{\prime} \cdot\left(h^{\prime}\right)^{2}$. Let $L=r^{\prime} \cdot\left(h^{\prime}\right)^{2}$, and consider instance $\operatorname{MBCS}(G, L)$ of the Maximum Bounded-Crossing Subgraph problem. From the above discussion, the value of the optimal solution to this problem is at least $|E(H)| \geq \frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{4 \log n}$. Therefore, by applying the $\alpha(N)$-approximation algorithm for Maximum Bounded-Crossing Subgraph to this instance, we obtain a solution $G^{\prime}$ to instance $\operatorname{MBCS}(G, L)$ of Maximum Bounded-Crossing Subgraph, whose value is at least $\frac{C^{*}}{4 \alpha(n) \log n}$. Since we have assumed that $C^{*} \geq 16 n \alpha(n) \log ^{7} n$, we get that $\left|E\left(G^{\prime}\right)\right| \geq 4 n \log ^{6} n$. We can now use the algorithm from Lemma 6.6 to compute integers $r^{\prime \prime}, h^{\prime \prime}>0$, such that $r^{\prime \prime} \cdot\left(h^{\prime \prime}\right)^{2} \leq L \cdot \log ^{6} n$, together with a good solution $\mathcal{H}^{\prime}$ to instance $\operatorname{GP}\left(G, r^{\prime \prime}, h^{\prime \prime}\right)$, whose value is at least $\Omega\left(\frac{\left|E\left(G^{\prime}\right)\right|}{\log n}\right) \geq \Omega\left(\frac{C^{*}}{\alpha(n) \log ^{2} n}\right)$. Notice however that it is possible that $r^{\prime \prime}>r^{\prime}$ or $h^{\prime \prime}>h^{\prime}$ hold, so the solution that we obtain may not be a valid solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning. We show that, if $\left(r^{\prime}, h^{\prime}\right)=\left(r^{*}, h^{*}\right)$, then $h^{\prime \prime}$ cannot be much larger than $h^{\prime}$. We then slightly modify solution $\mathcal{H}^{\prime}$, to transform it into a valid solution $\mathcal{H}_{r^{\prime}, h^{\prime}}$ to instance $\mathrm{GP}\left(G, r^{\prime}, h^{\prime}\right)$, while only decreasing the solution cost slightly. This completes the computation of the solution $\mathcal{H}_{r^{\prime}, h^{\prime}}$ associated with parameters $\left(r^{\prime}, h^{\prime}\right)$, and the algorithm for the first case.
Consider now the second case, where $C^{*}<16 n \alpha(n) \log ^{7} n$. In this case, we start by computing a
maximal subgraph $F$ of $G$, such that $F$ is a forest, with maximum vertex degree at most $h$. Let $S$ be the set of all vertices of $G$ that are adjacent to at least one edge of $F$, and denote $|S|=n^{\prime}$.
We consider two subcases of Case 2. The first subcase happens if $|E(F)| \geq \frac{C^{*}}{32 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}}$. In this case, we show an algorithm that decomposes $F$ into $r$ subgraphs containing at most $h$ edges each, so that the total number of edges in all such subgraphs is close to $|E(F)|$. Therefore, we obtain a solution $\mathcal{H}$ to instance $\operatorname{GP}(G, r, h)$, whose value is close to $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$. Consider now the second subcase, where $|E(F)|<\frac{C^{*}}{32 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}}$. We show that in this case, there is a solution to instance $\operatorname{GP}(G[S], r, h)$ of $(r, h)$-Graph Partitioning, whose value is at least $\frac{C^{*}}{2}$. From now on we only consider instane $\operatorname{GP}(G[S], r, h)$. We assume again that we are given the value $C^{* *}$ of the optimal solution to this instance, where $\frac{C^{*}}{2} \leq C^{* *} \leq C^{*}$. As before, this can be assumed since we can try all guesses for the value $C^{* *}$, and it is sufficient to ensure that the algorithm works correctly if the value $C^{* *}$ is guessed correctly. Recall that we have denoted $n^{\prime}=|S|$. Since $(h-1) \cdot|S| \leq|E(F)| \leq \frac{C^{*}}{32 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}}$, while $C^{* *} \geq \frac{C^{*}}{2}$, we get that $C^{* *} \geq 16 n^{\prime} \cdot \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}$.
We have now obtained a new instance $\operatorname{GP}(G[S], r, h)$ of ( $r, h$ )-Graph Partitioning, in which the value of the optimal solution $C^{* *} \geq 16 n \alpha(n) \log ^{7} n^{\prime}$, where $n^{\prime}=|S|$. We can now repeat our algorithm for Case 1, to obtain the desired approximate solution to instance $\operatorname{GP}(G[S], r, h)$, which, in turn will provide an approximate solution to the original instance $\operatorname{GP}(G, r, h)$.

We now turn to the formal proof of Theorem 6.2. We assume that we are given an instance $\operatorname{GP}(G, r, h)$ of the (r,h)-Graph Partitioning problem, where $|V(G)|=n$, together with a guess $C^{*}$ on the value $\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$ of the optimal solution to this instance. Our goal is to compute a solution $\mathcal{H}$ to instance $\operatorname{GP}(G, r, h)$ of $(r, h)$-Graph Partitioning, such that, if $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$, then the value of the solution $\mathcal{H}$ is at least $\Omega\left(\frac{C^{*}}{(\alpha(n))^{2} \operatorname{poly} \log n}\right)$. Note that we can assume that $n$ is greater than a sufficiently large constant, since otherwise we can solve the problem efficiently via exhaustive search. We distinguish between two cases, depending on whether $C^{*} \geq 16 n \alpha(n) \log ^{7} n$ holds.

### 6.3.1 Case 1: $C^{*} \geq 16 n \alpha(n) \log ^{7} n$

Applying Lemma 6.5 to the optimal solution tp instance $\operatorname{GP}(G, r, h)$, we conclude that there are positive integers $r^{*} \leq r, h^{*} \leq h$ and a collection $\mathcal{H}^{*}$ of subgraphs of $G$, such that $\mathcal{H}^{*}$ is a good solution to instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$, of $(r, h)$-Graph Partitioning. Moreover, if $C^{*}=\mathrm{OPT}_{\mathrm{GP}}(G, r, h)$, then $r^{*} \cdot h^{*} \geq \sum_{H \in \mathcal{H}^{*}}|E(H)| \geq \frac{C^{*}}{4 \log n}$. Our algorithm tries all possible values of integers $1 \leq r^{\prime} \leq r$ and $1 \leq h^{\prime} \leq h$ with $r^{\prime} h^{\prime} \geq \frac{C^{*}}{4 \log n}$. For each such pair $\left(r^{\prime}, h^{\prime}\right)$ of integers, we will compute a solution $\mathcal{H}_{r^{\prime}, h^{\prime}}$ to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$. At the end, our algorithm will output the best solution from among $\left\{\mathcal{H}_{r^{\prime}, h^{\prime}} \left\lvert\, r^{\prime} h^{\prime} \geq \frac{C^{*}}{4 \log n}\right.\right\}$. It is sufficient to ensure that, for integers $\left(r^{\prime}, h^{\prime}\right)=\left(r^{*}, h^{*}\right)$, the value of the resulting solution $\mathcal{H}_{r^{\prime}, h^{\prime}}$ is at least $\Omega\left(\frac{C^{*}}{(\alpha(n))^{2} \operatorname{poly} \log n}\right)$.
From now on we fix a pair $1 \leq r^{\prime} \leq r, 1 \leq h^{\prime} \leq h$ of integers, with $r^{\prime} h^{\prime} \geq \frac{C^{*}}{4 \log n}$. Let $L=r^{\prime} \cdot\left(h^{\prime}\right)^{2}$. We apply the $\alpha(N)$-approximation algorithm for the Maximum Bounded-Crossing Subgraph problem to instance $\operatorname{MBCS}(G, L)$, and obtain a solution that we denote by $H$. We use the following observation.

Observation 6.13. If $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h), r^{\prime}=r^{*}$ and $h^{\prime}=h^{*}$, then $|E(H)| \geq \frac{C^{*}}{4 \alpha(n) \log n}$ must hold.

Proof: Assume that $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h), r^{\prime}=r^{*}$ and $h^{\prime}=h^{*}$. Recall that there exists a good solution $\mathcal{H}^{*}$ to instance $\operatorname{GP}\left(G, r^{*}, h^{*}\right)$, with $\sum_{H^{\prime} \in \mathcal{H}^{*}}\left|E\left(H^{\prime}\right)\right| \geq \frac{C^{*}}{4 \log n}$. Consider the graph $\hat{H}=\bigcup_{H^{\prime} \in \mathcal{H}^{*}} H^{\prime}$. Since the crossing number of a graph $G^{\prime}$ may not be higher than $\left|E\left(G^{\prime}\right)\right|^{2}$, we get that $\operatorname{CrN}(\hat{H}) \leq \sum_{H^{\prime} \in \mathcal{H}^{*}} \operatorname{CrN}\left(H^{\prime}\right) \leq r^{\prime} \cdot\left(h^{\prime}\right)^{2}=L$. Therefore, $\hat{H}$ is a valid solution to instance $\operatorname{MBCS}(G, L)$, whose value is at least $\frac{C^{*}}{4 \log n}$. Since we use an $\alpha(N)$-approximation algorithm for Maximum Bounded-Crossing Subgraph, and $|V(G)|=n$, we get that $|E(H)| \geq \frac{C^{*}}{4 \alpha(n) \log n}$.
If $|E(H)|<\frac{C^{*}}{4 \alpha(n) \log n}$, then we terminate the algorithm and return an empty solution: in this case, we are guaranteed that either $C^{*}$, or $r^{\prime}, h^{\prime}$ are guessed incorrectly. Therefore, we assume from now on that $|E(H)| \geq \frac{C^{*}}{4 \alpha(n) \log n}$ holds. Note that, since in Case $1, C^{*} \geq 16 n \alpha(n) \log ^{7} n$ holds, we are guaranteed that $|E(H)| \geq 4 n \log ^{6} n$.

Next, we apply the algorithm from Lemma 6.6 to instance $\operatorname{MBCS}(G, L)$ of Maximum BoundedCrossing Subgraph, to compute integers $r^{\prime \prime}, h^{\prime \prime}>0$, such that $r^{\prime \prime} \cdot\left(h^{\prime \prime}\right)^{2} \leq L \cdot \log ^{6} n=r^{\prime} \cdot\left(h^{\prime}\right)^{2} \cdot \log ^{6} n$, and $r^{\prime \prime} \cdot h^{\prime \prime} \geq \Omega\left(\frac{|E(H)|}{\log n}\right) \geq \Omega\left(\frac{C^{*}}{\alpha(n) \log ^{2} n}\right)$. The algorithm also computes a good solution $\mathcal{H}^{\prime}$ to instance $\operatorname{GP}\left(G, r^{\prime \prime}, h^{\prime \prime}\right)$ of $(r, h)$-Graph Partitioning. Note that, while the value of the solution $\mathcal{H}^{\prime}$ to instance $\operatorname{GP}\left(G, r^{\prime \prime}, h^{\prime \prime}\right)$ is guaranteed to be close to $C^{*}$, we are only guaranted that $\mathcal{H}^{\prime}$ is a valid solution to instance $\operatorname{GP}\left(G, r^{\prime \prime}, h^{\prime \prime}\right)$ of the problem, and it may not be a valid solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$. In our next steps, we will either correctly established that at least one of $C^{*}, r^{\prime}, h^{\prime}$ was not guessed correctly; or we will slightly modify $\mathcal{H}^{\prime}$ to obtain a valid solution to intance $\mathrm{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning, whose value remains close to that of $\mathcal{H}^{\prime}$. We start with the following observation.

Observation 6.14. There is a large enough constant $c^{\prime}$, such that, if $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h), r^{\prime}=r^{*}$ and $h^{\prime}=h^{*}$ hold, then $h^{\prime \prime} \leq c^{\prime} h^{\prime} \cdot \alpha(n) \cdot \log ^{8} n$.

Proof: Recall that we have established that:

$$
r^{\prime \prime} \cdot h^{\prime \prime} \geq \Omega\left(\frac{C^{*}}{\alpha(n) \log ^{2} n}\right)
$$

Assume for contradicton that $h^{\prime \prime}>c^{\prime} h^{\prime} \cdot \alpha(n) \cdot \log ^{8} n$, where $c^{\prime}$ is a large enough constant. Then:

$$
r^{\prime \prime} \cdot\left(h^{\prime \prime}\right)^{2} \geq 4 h^{\prime} \cdot C^{*} \log ^{6} n .
$$

Notice that $C^{*} \geq r^{\prime} h^{\prime} / 2$ must hold. Indeed, since $r^{\prime} \leq r$ and $h^{\prime} \leq h$, any solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of $(r, h)$-Graph Partitioning is also a feasible solution to instance $\operatorname{GP}(G, r, h)$. Since we have assumed that $r^{\prime}=r^{*}$ and $h^{\prime}=h^{*}$, there is a good solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of $(r, h)$-Graph Partitioning, and the value of any such good solution is at least $r^{\prime} h^{\prime} / 2$. Since we have assumed that $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$, we get that $r^{\prime} h^{\prime} / 2 \leq C^{*}$ must hold. We conclude that, if $h^{\prime \prime}>h^{\prime} \cdot \alpha(n) \cdot \log ^{8} n$, then:

$$
r^{\prime \prime} \cdot\left(h^{\prime \prime}\right)^{2} \geq 2 r^{\prime}\left(h^{\prime}\right)^{2} \log ^{6} n
$$

But we have already established above that $r^{\prime \prime} \cdot\left(h^{\prime \prime}\right)^{2} \leq r^{\prime} \cdot\left(h^{\prime}\right)^{2} \cdot \log ^{6} n$, a contradiction.

We will now slightly modify the collection $\mathcal{H}^{\prime}$ of subgraphs of $G$ to obtain a feasible solution to instance $\operatorname{GP}(G, r, h)$, whose value is close to the value of $\mathcal{H}^{\prime}$. First, for every cluster $H \in \mathcal{H}^{\prime}$, if $|E(H)|>h^{\prime} / 2$, then we discard arbitrary edges from graph $H$, until $|E(H)|=h^{\prime} / 2$ holds. From Observation 6.14, if $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h), r^{\prime}=r^{*}$ and $h^{\prime}=h^{*}$ hold, then the total number of edges in the graphs of $\mathcal{H}^{\prime}$ decreases by at most factor $O\left(\alpha(n) \cdot \log ^{8} n\right)$ as the result of this transformation, and so $\sum_{H \in \mathcal{H}^{\prime}}|E(H)| \geq \Omega\left(\frac{r^{\prime \prime} \cdot h^{\prime \prime}}{\alpha(n) \cdot \log ^{8} n}\right) \geq \Omega\left(\frac{C^{*}}{(\alpha(n))^{2} \log ^{10} n}\right)$ holds. Also, if, at the end of this transformation, $\sum_{H \in \mathcal{H}^{\prime}}|E(H)|>\frac{h^{\prime} \cdot r^{\prime}}{2}$ holds, then we discard arbitrary edges from the graphs in $\mathcal{H}^{\prime}$ until $\sum_{H \in \mathcal{H}^{\prime}}|E(H)| \leq \frac{h^{\prime} \cdot r^{\prime}}{2}$ holds. Since $h^{\prime} r^{\prime} \geq \frac{C^{*}}{4 \log n}, \sum_{H \in \mathcal{H}^{\prime}}|E(H)| \geq \Omega\left(\frac{C^{*}}{(\alpha(n))^{2} \log ^{10} n}\right)$ continues to hold.
If $\left|\mathcal{H}^{\prime}\right| \leq r^{\prime}$, then we have obtained a valid solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of value $\Omega\left(\frac{C^{*}}{(\alpha(n))^{2} \log ^{10} n}\right)$. Otherwise, we perform further transformations to the set $\mathcal{H}^{\prime}$ of graphs as follows.

While $\left|\mathcal{H}^{\prime}\right|>r^{\prime}$, we let $H^{\prime}, H^{\prime \prime} \in \mathcal{H}^{\prime}$ be a pair of graphs with smallest number of edges, breaking ties arbitrarily. We remove $H^{\prime}$ and $H^{\prime \prime}$ from $\mathcal{H}^{\prime}$, and we add a new graph $H=H^{\prime} \cup H^{\prime \prime}$ to $\mathcal{H}^{\prime}$ instead. The procedure is terminated once $\left|\mathcal{H}^{\prime}\right|=r^{\prime}$ holds. We claim that at the end of this procedure, for every graph $H \in \mathcal{H}^{\prime},|E(H)| \leq h^{\prime}$ holds. Indeed, assume otherwise. Consider the first time when a graph $H$ with $|E(H)|>h^{\prime}$ was added to $\mathcal{H}^{\prime}$. Then $H=H^{\prime} \cup H^{\prime \prime}$ must hold, where $H^{\prime}, H^{\prime \prime}$ are two graphs that belonged to $\mathcal{H}^{\prime}$ prior to this iterations. Then at least one of these two graphs must contain more than $h^{\prime} / 2$ edges. From the choice of the graphs $H^{\prime}, H^{\prime \prime}$, and from the fact that $\left|\mathcal{H}^{\prime}\right|>r^{\prime}$ held at the beginning of the iteration, we get that, at the beginning of the iteration, there were at least $r^{\prime}$ graphs $\tilde{H} \in \mathcal{H}^{\prime}$ with $|E(\tilde{H})|>h^{\prime} / 2$. But then $\sum_{\tilde{H} \in \mathcal{H}^{\prime}}|E(\tilde{H})|>\frac{r^{\prime} h^{\prime}}{2}$ held at the beginning of the iteration. Since the total number of edges contained in the graphs of $\mathcal{H}^{\prime}$ does not change over the course of the algorithm, we reach a contradiction, since we have ensured that, at the beginning of the algorithm, $\sum_{\tilde{H} \in \mathcal{H}^{\prime}}|E(\tilde{H})| \leq \frac{h^{\prime} \cdot r^{\prime}}{2}$ held. We return the resulting collection $\mathcal{H}^{\prime}$ of subgraphs of $G$, which is guaranteed to be a feasible solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$, of value at least $\Omega\left(\frac{C^{*}}{(\alpha(n))^{2} \log ^{10} n}\right)$.

### 6.3.2 Case 2: $C^{*}<16 n \alpha(n) \log ^{7} n$

In this case, we start by computing a maximal subgraph $F$ of $G$, such that $F$ is a forest, and maximum vertex degree in $F$ is at most $h$. Such a graph $F$ can be computed via a simple greedy algorithm. We start with graph $F$ containing the set $V(F)=V(G)$ of vertices and no edges. We then consider the edges of $G$ one by one. For each such edge $e \in E(G)$, if graph $F \cup\{e\}$ remains a forest with maximum vertex degree at most $h$, then we add $e$ to $F$. Once every edge of $G$ is processed, we obtain the final graph $F$. Let $S$ be the set of all vertices of $G$ that are adjacent to at least one edge of $F$, and denote $|S|=n^{\prime}$.
We consider two subcases of Case 2. The first subcase, Case 2a happens if $|E(F)| \geq \frac{C^{*}}{64 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}}$. In this case, we use the following simple observation, that will allow us to decompose the forest $F$ to obtain a solution to instance $\operatorname{GP}\left(G, r^{\prime}, h^{\prime}\right)$ of (r,h)-Graph Partitioning, whose value is close to $|E(F)|$.

Observation 6.15. There is an efficient algorithm, that, given a tree $T$ with maximum vertex degree at most $\Delta$ and $|E(T)| \geq \Delta / 2$, computes a collection $\mathcal{T}$ of vertex-disjoint subgraphs of $T$,
such that for each subgraph $T^{\prime} \in \mathcal{T}, \frac{\Delta}{2} \leq\left|E\left(T^{\prime}\right)\right| \leq \Delta$, and $\sum_{T^{\prime} \in \mathcal{T}}\left|E\left(T^{\prime}\right)\right| \geq \frac{|E(T)|}{2}$.

Proof: We root the tree $T$ in an arbitrary vertex $v$. Initially, we let $\mathcal{T}=\emptyset$. As long as $|E(T)|>\Delta$, we perform iterations. In every iteration, we consider an arbitrary vertex $u$ in the current tree $T$, that is a non-leaf vertex, but all children of $u$ are leaf vertices. Let $T^{\prime}$ be the subtree of $T$ rooted at vertex $u$. Note that $1 \leq\left|E\left(T^{\prime}\right)\right| \leq \Delta-1$. We add graph $T^{\prime}$ to $\mathcal{T}$, and we delete all vertices of $T^{\prime}$ from $T$. Note that, as the result of this iteration, the unique edge $e$ connecting $u$ to its parent-vertex in the tree $T$ is deleted from $T$, and it does not belong to any graph in $\mathcal{T}$. We let $e^{\prime} \in E(T)$ be an arbitrary edge (which must exist since $|E(T)| \geq 1$ ), and we say that $e^{\prime}$ is responsible for the deletion of the edge $e$. We then continue to the next iteration. The algorithm terminates once $|E(T)| \leq \Delta$ holds. We then add graph $T$ to the collection $\mathcal{T}$ and terminate the algorithm. It is immediate to verify that, for every graph $T^{\prime} \in \mathcal{T},\left|E\left(T^{\prime}\right)\right| \leq \Delta$ holds. Moreover, if an edge $e$ belonged to the original graph $T$, and it does not belong to $\bigcup_{T^{\prime} \in \mathcal{T}} E\left(T^{\prime}\right)$, then some edge of $\bigcup_{T^{\prime} \in \mathcal{T}} E\left(T^{\prime}\right)$ is designated as being responsible for deleting $e$. It is easy to verify that every edge $e^{\prime} \in \bigcup_{T^{\prime} \in \mathcal{T}} E\left(T^{\prime}\right)$ may be responsible for the deletion of at most one edge. Therefore, at the end of the algorithm, $\sum_{T^{\prime} \in \mathcal{T}}\left|E\left(T^{\prime}\right)\right| \geq|E(T)| / 2$ holds.
Initially, we construct a collection $\mathcal{H}^{\prime}$ of subgraphs of $F$ as follows. For every tree $T$ of the forest $F$, if $|V(T)| \leq h$, then we add $T$ to $\mathcal{H}^{\prime}$. Otherwise, we apply Observation 6.15 to tree $T$ with parameter $\Delta=h$, and add the graphs in the resulting collection $\mathcal{T}$ to $\mathcal{H}^{\prime}$. At the end of this algorithm, for every graph $H^{\prime} \in \mathcal{H}^{\prime},\left|E\left(H^{\prime}\right)\right| \leq h$ holds, and $\sum_{H^{\prime} \in \mathcal{H}^{\prime}}\left|E\left(H^{\prime}\right)\right| \geq \frac{|E(F)|}{2} \geq \frac{C^{*}}{128 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}}$. For every graph $H^{\prime} \in \mathcal{H}^{\prime}$, if $\left|E\left(H^{\prime}\right)\right|>h / 2$, then we delete edges from $H^{\prime}$ until $\left|E\left(H^{\prime}\right)\right| \leq$ $h / 2$ holds. Clearly, after this transformation, $\sum_{H^{\prime} \in \mathcal{H}^{\prime}}\left|E\left(H^{\prime}\right)\right| \geq \frac{C^{*}}{256 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}} \geq \Omega\left(\frac{C^{*}}{\alpha(n) \cdot \log ^{7} n}\right)$ holds. If $\sum_{H \in \mathcal{H}^{\prime}}|E(H)|>h r / 2$, then we discard arbitrary edges from the graphs in $\mathcal{H}^{\prime}$, until $\sum_{H \in \mathcal{H}^{\prime}}|E(H)|=h r / 2$ holds. Since, if $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h), C^{*} \leq r h$, we are still guaranteed that $\sum_{H^{\prime} \in \mathcal{H}^{\prime}}\left|E\left(H^{\prime}\right)\right| \geq \Omega\left(\frac{C^{*}}{\alpha(n) \cdot \log ^{7} n}\right)$ holds.
If $\left|\mathcal{H}^{\prime}\right| \leq r$, then we have obtained a valid solution to instance $\mathrm{GP}(G, r, h)$ of value $\Omega\left(\frac{C^{*}}{\alpha(n) \cdot \log ^{7} n}\right)$. Otherwise, we proceed exactly like in Case 1 in order to transform $\mathcal{H}^{\prime}$ into a valid solution to instance $\operatorname{GP}(G, r, h)$ of ( $r, h$ )-Graph Partitioning, without changing the total number of edges that lie in the graphs of $\mathcal{H}^{\prime}$. While $\left|\mathcal{H}^{\prime}\right|>r$, we let $H^{\prime}, H^{\prime \prime} \in \mathcal{H}^{\prime}$ be a pair of graphs with smallest number of edges, breaking ties arbitrarily. We remove $H^{\prime}$ and $H^{\prime \prime}$ from $\mathcal{H}^{\prime}$, and we add a new graph $H=H^{\prime} \cup H^{\prime \prime}$ to $\mathcal{H}^{\prime}$ instead. The procedure is terminated once $\left|\mathcal{H}^{\prime}\right|=r$ holds. We claim that at the end of this procedure, for every graph $H \in \mathcal{H}^{\prime},|E(H)| \leq h$ holds. Indeed, assume otherwise. Consider the first time when a graph $H$ with $|E(H)|>h$ was added to $\mathcal{H}^{\prime}$. Then $H=H^{\prime} \cup H^{\prime \prime}$ must hold, where $H^{\prime}, H^{\prime \prime}$ are two graphs that belonged to $\mathcal{H}^{\prime}$ prior to this iterations. Then at least one of these two graphs must contain more than $h / 2$ edges. From the choice of the graphs $H^{\prime}, H^{\prime \prime}$, and from the fact that $\left|\mathcal{H}^{\prime}\right|>r$ held at the beginning of the iteration, we get that, at the beginning of the iteration, there were at least $r$ graphs $\tilde{H} \in \mathcal{H}^{\prime}$ with $|E(\tilde{H})|>h / 2$. But then $\sum_{\tilde{H} \in \mathcal{H}^{\prime}}|E(\tilde{H})|>\frac{r h}{2}$ held at the beginning of the iteration. Since the total number of edges contained in the graphs of $\mathcal{H}^{\prime}$ does not change over the course of the algorithm, we reach a contradiction, since we have ensured that, at the beginning of the algorithm, $\sum_{\tilde{H} \in \mathcal{H}^{\prime}}|E(\tilde{H})| \leq \frac{h \cdot r}{2}$ held. We return the resulting set $\mathcal{H}^{\prime}$ of subgraphs of $G$, which is guaranteed to be a feasible solution to instance $\operatorname{GP}(G, r, h)$ of $(r, h)$-Graph Partitioning, of value at least $\Omega\left(\frac{C^{*}}{\alpha(n) \log ^{7} n}\right)$.

It now remains to consider Case (2b), where $C^{*}<16 n \alpha(n) \log ^{7} n$ and $|E(F)|<\frac{C^{*}}{64 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}}$. Recall that $S$ is the set of all vertices of $G$ that are adjacent to at least one edge of $F$, and recall that we have denoted $|S|=n^{\prime}$.
In this case, we let $G^{\prime}=G[S]$, and we consider instance $\operatorname{GP}\left(G^{\prime}, r, h\right)$ of ( $r, h$ )-Graph Partitioning. Notice that, if $\mathcal{H}^{\prime}$ is a valid solution to instance $\operatorname{GP}\left(G^{\prime}, r, h\right)$ of ( $r, h$ )-Graph Partitioning, then it is also a valid solution to instance $\operatorname{GP}(G, r, h)$ of $(r, h)$-Graph Partitioning. We start by showing that $\operatorname{OPT}_{\mathrm{GP}}\left(G^{\prime}, r, h\right)$ is close to $C^{*}$.
Observation 6.16. If $C^{*}=\operatorname{OPT}_{\mathrm{GP}}(G, r, h)$, and Case (2b) happens, then $\mathrm{OPT}_{\mathrm{GP}}\left(G^{\prime}, r, h\right) \geq \frac{C^{*}}{2}$.
Proof: Let $S^{\prime} \subseteq S$ be the set of all vertices whose degree in $F$ is $h$, and let $S^{*}$ the set of all vertices of $F$ that are isolated.

Let $\mathcal{H}$ be the optimal solution to instance $\operatorname{GP}(G, r, h)$, and let $E^{\prime}=\bigcup_{H \in \mathcal{H}} E(H)$. We partition the set $E^{\prime}$ of edges into two subsets: set $E_{1}^{\prime}$ containing all edges that lie in $G^{\prime}=G[S]$, and set $E_{2}^{\prime}$ containing all remaining edges. Clearly, for every edge $e=(x, y) \in E_{2}^{\prime}$, at least one endpoint of $e$ must lie in $S^{*}$. Assume w.l.o.g. that $x \in S^{*}$. We claim that $y \in S^{\prime}$ must hold. Indeed, otherwise $F \cup\{e\}$ remains a forest, with maximum vertex degree at most $h$, contradicting the fact that $F$ is a maximal subgraph of $G$ with these properties.
Therefore, very edge of $E_{2}^{\prime}$ connects a vertex of $S^{*}$ to a vertex of $S^{\prime}$. We claim that $\left|E_{2}^{\prime}\right| \leq\left|S^{\prime}\right| \cdot h$. Indeed, from the definition of the (r,h)-Graph Partitioning problem, for every graph $H \in \mathcal{H},|E(H)| \leq$ $h$, and all graphs in $\mathcal{H}$ are disjoint in their vertices. Therefore, every vertex of $G$ may be incident to at most $h$ edges of $E^{\prime}$. Since every edge of $E_{2}^{\prime}$ has a vertex of $S^{\prime}$ as its endpoint, we get that $\left|E_{2}^{\prime}\right| \leq\left|S^{\prime}\right| \cdot h$.
Recall that, from our definition, every vertex $v \in S^{\prime}$ has degree $h$ in $F$. Therefore, $|E(F)| \geq$ $(h-1) \cdot\left|S^{\prime}\right|$. We conclude that $|E(F)| \geq\left|E_{2}^{\prime}\right| / 2$, and so $\left|E_{2}^{\prime}\right| \leq 2|E(F)|<\frac{C^{*}}{4}$. Since $\left|E^{\prime}\right|=C^{*}$, we get that $\left|E_{1}^{\prime}\right| \geq C^{*} / 2$.

We now define a solution $\mathcal{H}^{\prime}$ to instance $\operatorname{GP}\left(G^{\prime}, h, r\right)$ of (r,h)-Graph Partitioning. For every graph $H \in \mathcal{H}$, we let $H^{\prime}$ be a graph that is obtained from $H$ by deleting all vertices of $S^{*}$ from it, and we let $\mathcal{H}^{\prime}=\left\{H^{\prime} \mid H \in \mathcal{H}\right\}$. It is easy to verify that $\mathcal{H}^{\prime}$ is a valid solution to instance $\operatorname{GP}\left(G^{\prime}, r, h\right)$, and that its value is at least $\left|E^{\prime} \backslash E_{2}^{\prime}\right| \geq \frac{C^{*}}{2}$. We conclude that $\mathrm{OPT}_{\mathrm{GP}}\left(G^{\prime}, r, h\right) \geq \frac{C^{*}}{2}$.
Denote $C^{\prime}=\operatorname{OPT}_{\mathrm{GP}}\left(G^{\prime}, r, h\right)$. From the above discussion $\frac{C^{*}}{2} \leq C^{\prime} \leq C^{*}$. Notice that for every tree $T$ of $F$, if $T$ is not a singleton vertex, then $|E(T)| \geq|V(T)|-1 \geq \frac{|V(T)|}{2}$. Therefore, $|E(F)| \geq$ $\frac{|S|}{2}=\frac{n^{\prime}}{2}$. On the other hand, from our assumpution, $|E(F)|<\frac{C^{*} \cdot 2}{64 \alpha\left(n^{\prime}\right) \cdot \log ^{7} n^{\prime}}$. We then conclude that $C^{*}>32 n^{\prime} \alpha\left(n^{\prime}\right) \cdot \log ^{7} n$.
We will now focus on solving instance $\operatorname{GP}\left(G^{\prime}, r, h\right)$ of the ( $r, h$ )-Graph Partitioning problem. As before, we will try all guesses $C^{* *}$ on the value $C^{\prime}$ of the optimal solution for this problem. Note that we only need to consider values $C^{* *}$ that are integers, with $\frac{C^{*}}{2} \leq C^{* *} \leq C^{*}$. Furthermore, from the above discussion, for each such guess, $C^{* *} \geq \frac{C^{*}}{2} \geq 16 n^{\prime} \alpha\left(n^{\prime}\right) \cdot \log ^{7} n$ holds, so Case 1 will occur. We execute the algorithm from Case 1 for each such guessed value $C^{* *}$ and output the best among the resulting solutions. We are then guaranteed to obtain a solution $\mathcal{H}$ to instance $\operatorname{GP}\left(G^{\prime}, r, h\right)$ of value at least $\Omega\left(\frac{C^{\prime}}{(\alpha(n))^{2} \log ^{10} n}\right) \geq \Omega\left(\frac{C^{*}}{(\alpha(n))^{2} \log ^{10} n}\right)$. Clearly, $\mathcal{H}$ is also a valid solution to instance $\operatorname{GP}(G, r, h)$. Overall, we obtain an efficient $O\left((\alpha(n))^{2} \cdot\right.$ poly $\left.\log n\right)$-approximation algorithm for
(r,h)-Graph Partitioning.

## 7 Acknowledgement

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

## A Proof of Lemma 2.1

We prove each of the directions of the reductions separately, in the following two subsections.

## A. 1 Reduction from Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph to Densest $k$-Subgraph

Assume that exists an $\alpha(n)$-approximation algorithm $\mathcal{A}$ for the Densest $k$-Subgraph problem with running time at most $T(n)$, where $n$ is the number of vertices in the input graph. We show an $O\left(\alpha\left(N^{2}\right)\right)$-approximation algorithm for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, whose running time is at most $O\left(T\left(N^{2}\right) \cdot \operatorname{poly}(N)\right)$, where $N$ is the number of vertices in the input graph.

Let $\operatorname{DkS}\left(G, k_{1}, k_{2}\right)$ be the input instance to the Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph problem. Denote $G=(A, B, E)$, so $|A \cup B|=N$. We construct another bipartite graph $H=\left(A^{\prime}, B^{\prime}, E^{\prime}\right)$, that will serve as input to the Densest $k$-Subgraph problem, as follows. We define, for every vertex $u \in A$, a collection $T_{u}=\left\{u^{1}, \ldots, u^{k_{2}}\right\}$ of vertices that we call copies of $u$, and we let $A^{\prime}=\bigcup_{u \in A} T_{u}$. Similarly, we define, for every vertex $v \in B$, a set $T_{v}=\left\{v^{1}, \ldots, v^{k_{1}}\right\}$ of $k_{1}$ vertices, that we call copies of $v$, and we let $B^{\prime}=\bigcup_{v \in B} T_{v}$. The set $E^{\prime}$ of edges of $H$ contains, for every edge $e=(u, v) \in E(G)$, all edges in $T_{u} \times T_{v}$. Note that $|V(H)| \leq \max \left\{k_{1}, k_{2}\right\} \cdot|A \cup B| \leq N^{2}$.

Let $k=2 k_{1} k_{2}$, and consider the instance $\operatorname{DkS}(H, k)$ of the Densest $k$-Subgraph problem. We use the following observation to lower-bound its optimal solution cost.
Observation A.1. $\mathrm{OPT}_{\mathrm{DkS}}(H, k) \geq k_{1} k_{2} \cdot \mathrm{OPT}_{\mathrm{BDkS}}\left(G, k_{1}, k_{2}\right)$.
Proof: Let $S^{*}$ be the optimal solution to instance $\operatorname{BDkS}\left(G, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$ Subgraph. Denote $S_{A}^{*}=S^{*} \cap A$ and $S_{B}^{*}=S^{*} \cap B$, so $\left|S_{A}^{*}\right|=k_{1}$ and $\left|S_{B}^{*}\right|=k_{2}$ hold. We define $T_{A}^{*}=\bigcup_{u \in S_{A}^{*}} T_{u}$ and $T_{B}^{*}=\bigcup_{v \in S_{B}^{*}} T_{v}$. From the construction of $H$, it is clear that $\left|T_{A}^{*}\right|=\left|T_{B}^{*}\right|=$ $k_{1} \cdot k_{2}$, and $\left|E_{H}\left(T_{A}^{*}, T_{B}^{*}\right)\right|=k_{1} k_{2} \cdot\left|E_{G}\left(S_{A}^{*}, S_{B}^{*}\right)\right|$. Therefore, $T_{A}^{*} \cup T_{B}^{*}$ is a feasible solution to instance $\operatorname{DkS}(H, k)$ of Densest $k$-Subgraph, and so $\mathrm{OPT}_{\mathrm{DkS}}\left(H, 2 k_{1} k_{2}\right) \geq k_{1} k_{2} \cdot \mathrm{OPT}_{\mathrm{BDkS}}\left(G, k_{1}, k_{2}\right)$.
In order to complete the reduction, we need the following claim.
Claim A.2. There is an efficient algorithm, that, given any solution $W$ to the instance $\operatorname{DkS}(H, k)$ of the Densest $k$-Subgraph problem, computes a solution to instance $\operatorname{BDkS}\left(G, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph, whose value is at least $\left|E_{H}(W)\right| /\left(4 k_{1} k_{2}\right)$.

Proof: Denote $W_{A}=W \cap A$ and $W_{B}=W \cap B$. We start by computing a partition $\left(W_{A}^{0}, \ldots, W_{A}^{2 k_{2}-1}\right)$ of the vertices of $W_{A}$ into $2 k_{2}$ subsets, containing at most $k_{1}$ vertices each, so that, for every vertex $u$ of $A$, no two copies of $u$ appear in the same subset.
In order to do so, we let $\sigma$ be an arbitrary ordering of the vertices of $W_{A}$, in which, for every vertex $u \in A$, all copies of $u$ that belong to $W_{A}$ appear consecutively. For all $0 \leq i \leq 2 k_{2}-1$, we let $W_{A}^{i} \subseteq W_{A}$ to be the set of all vertices $x \in W_{A}$, whose index is $i\left(\bmod 2 k_{2}\right)$ in this ordering. Since, for every vertex $u \in A,\left|T_{u}\right|=k_{2}$, it is immediate to verify that all copies of $u$ in $W_{A}$ lie in distinct sets. It is also immediate to verify that, for all $0 \leq i<2 k_{2},\left|W_{A}^{i}\right| \leq k_{1}$.
We similarly compute a partition $\left(W_{B}^{0}, \ldots, W_{B}^{2 k_{1}-1}\right)$ of the vertices of $W_{B}$ into $2 k_{1}$ subsets, containing at most $k_{2}$ vertices each, so that, for every vertex $v$ of $B$, no two copies of $v$ appear in the same subset.

Let $0 \leq i^{*}<2 k_{2}, 0 \leq j^{*}<2 k_{1}$ be a pair of indices, for which $\left|E_{H}\left(W_{A}^{i^{*}}, W_{B}^{j^{*}}\right)\right|$ is maximized. Clearly, $\left|E_{H}\left(W_{A}^{i^{*}}, W_{B}^{j^{*}}\right)\right| \geq\left|E_{H}(W)\right| /\left(4 k_{1} k_{2}\right)$. Finally, let $X \subseteq V(G)$ be the set of vertices containing every vertex $u \in V(G)$, whose copy lies in $W_{A}^{i^{*}} \cup W_{B}^{j^{*}}$. Note that $|X \cap A|=\left|W_{A}^{i^{*}}\right| \leq k_{1}$ and $|X \cap B|=$ $\left|W_{B}^{j^{*}}\right| \leq k_{2}$ must hold, so $X$ is a valid solution to instance $\operatorname{BDkS}\left(G, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph. Since $W_{A}^{i^{*}} \cup W_{B}^{j^{*}}$ contains at most one copy of every vertex of $V(G)$, from the construction of graph $H$, it is easy to verify that $\left|E_{G}(X)\right|=\left|E_{H}\left(W_{A}^{i^{*}}, W_{B}^{j^{*}}\right)\right| \geq\left|E_{H}(W)\right| /\left(4 k_{1} k_{2}\right)$.

We are now ready to complete our reduction. We apply the approximation algorithm $\mathcal{A}$ for the Densest $k$-Subgraph problem to instance $\operatorname{DkS}(H, k)$, to obtain a solution $W$. Since $|V(H)| \leq N^{2}$, and since $\mathcal{A}$ is a factor- $\alpha(n)$ approximation algorithm, from Observation A.1, we get that:

$$
\left|E_{H}(W)\right| \geq \frac{\mathrm{OPT}_{\mathrm{DkS}}(H, k)}{\alpha\left(N^{2}\right)} \geq \frac{k_{1} \cdot k_{2} \cdot \mathrm{OPT}_{\mathrm{BDkS}}\left(G, k_{1}, k_{2}\right)}{\alpha\left(N^{2}\right)}
$$

Additionally, the running time of the algorithm is $O\left(T\left(N^{2}\right)\right)$.
We then apply the algorithm from Claim A.2. whose running time is bounded by $O(\operatorname{poly}(N))$ to solution $W$ to instance $\operatorname{DkS}(H, k)$, to obtain a solution $X$ to instance $\operatorname{DkS}\left(G, k_{1}, k_{2}\right)$ of Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph. We are guaranteed that:

$$
\left|E_{G}(X)\right| \geq \frac{\left|E_{H}(W)\right|}{4 k_{1} k_{2}} \geq \Omega\left(\frac{\operatorname{OPT}_{\mathrm{DkS}}\left(G, k_{1}, k_{2}\right)}{\alpha\left(N^{2}\right)}\right)
$$

It is easy to verify that the running time of the algorithm is bounded by $O\left(T\left(N^{2}\right) \cdot \operatorname{poly}(N)\right)$.

## A. 2 Reduction from Densest $k$-Subgraph to Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph

We now assume that there exists an efficient $\alpha(N)$-approximation algorithm $\mathcal{A}^{\prime}$ for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, where $N$ is the number of vertices in the input graph. We show that there exists an efficient $O(\alpha(2 n))$-approximation algorithm for the Densest $k$-Subgraph problem, where $n$ is the number of vertices in the input graph.
Let $\operatorname{DkS}(G, k)$ be the input instance for the Densest $k$-Subgraph problem, so $|V(G)|=n$. We construct a bipartite graph $H=\left(V_{1}, V_{2}, E\right)$, where the vertex sets are $V_{1}=\left\{u^{1} \mid u \in V(G)\right\}$, $V_{2}=\left\{u^{2} \mid u \in V(G)\right\}$, and the edge set is $E=\left\{\left(u^{1}, v^{2}\right),\left(u^{2}, v^{1}\right) \mid(u, v) \in E(G)\right\}$. We denote $N=|V(H)|=2 n$. Consider the instance $\operatorname{DkS}\left(H, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph, where $k_{1}=k_{2}=k$. We use the following observation to lower-bound the optimal solution cost of this instance.

Observation A.3. $\operatorname{OPT}_{\mathrm{BDkS}}\left(H, k_{1}, k_{2}\right) \geq 2 \cdot \operatorname{OPT}_{\mathrm{DkS}}(G, k)$.

Proof: Let $V^{*}$ be the optimal solution to instance $\operatorname{DkS}(G, k)$ of Densest $k$-Subgraph. We define $U^{*}=\left\{v^{1}, v^{2} \mid v \in V^{*}\right\}$, so $\left|U^{*} \cap V_{1}\right|=k$ and $\left|U^{*} \cap V_{2}\right|=k$. Clearly, $U^{*}$ is a feasible solution to instance $\operatorname{BDkS}(H, k, k)$. Moreover, it is easy to verify that $\left|E_{H}\left(U^{*}\right)\right|=2 \cdot\left|E_{G}\left(V^{*}\right)\right|$, so $\operatorname{OPT}_{\mathrm{BDkS}}\left(H, k_{1}, k_{2}\right) \geq 2 \cdot \mathrm{OPT}_{\mathrm{DkS}}(G, k)$.

We apply Algorithm $\mathcal{A}^{\prime}$ to instance, $\operatorname{DkS}\left(H, k_{1}, k_{2}\right)$ of Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph, obtaining a solution $U^{\prime}$. We denote $U_{1}=U^{\prime} \cap V_{1}$ and $U_{2}=U^{\prime} \cap V_{2}$, so $\left|U_{1}\right|=\left|U_{2}\right|=k$, and $\left|E_{H}\left(U_{1}, U_{2}\right)\right| \geq$ $\operatorname{OPT}_{\mathrm{BDkS}}(H, k, k) / \alpha(N) \geq 2 \mathrm{OPT}_{\mathrm{DkS}}(G, k) / \alpha(2 n)$ from Observation A.3.
Let $U=\left\{v \mid v^{1} \in U_{1}\right.$ or $\left.v^{2} \in U_{2}\right\}$ be a subset of vertices of $G$. Clearly, $|U| \leq 2 k$, and $\left|E_{G}(U)\right| \geq$ $\frac{\left|E_{H}\left(U^{*}\right)\right|}{2} \geq \frac{\operatorname{OPT}_{\mathrm{DKS}}(G, k)}{\alpha(2 n)}$.
We then apply the algorithm from Lemma 2.3 to graph $G$ and set $U$ of vertices, with parameter $\beta=1 / 2$, to obtain a set $\tilde{U} \subseteq U$ of vertices, with $|\tilde{U}| \leq k$, and $\left|E_{G}(\tilde{U})\right| \geq \Omega\left(\left|E_{G}(U)\right|\right) \geq$ $\Omega\left(\operatorname{OPT}_{\mathrm{DkS}}(G, k) / \alpha(2 n)\right)$. Therefore, we obtained an $O(\alpha(2 n))$-approximate solution to instance $\operatorname{DkS}(G, k)$ of the Densest $k$-Subgraph problem.

## B Reduction from (r,h)-Graph Partitioning to Densest $k$-Subgraph

In this section we complete the proof Theorem 4.1, by showing a reduction from (r,h)-Graph Partitioning to Densest $k$-Subgraph. The reduction is very similar to the reduction from Dense $k$-Coloring to Densest $k$-Subgraph described in Section 4.

We start by formulating an LP-relaxation of the problem, whose number of constraints is bounded by $O(N)$, but the number of variables may be large. We then show an LP-rounding algorithm for this LP-relaxation, whose running time is $O(\operatorname{poly}(N))$ if it is given a solution to the LP-relaxation whose support size is bounded by $O(\operatorname{poly}(N))$. In order to compute an approximate LP-solution whose support size is sufficiently small, we design an approximate separation oracle for the dual of the LP-relaxation. We start with describing the LP-relaxation and providing an LP-rounding algorithm for it.

## B. 1 Linear Programming Relaxation and an LP-Rounding Algortihm

Let $\operatorname{GP}(G, r, h)$ be the input instance of $(r, \mathrm{~h})$-Graph Partitioning, and denote $|V(G)|=N$. We let $\mathcal{H}$ be the collection of all subgraphs $H \subseteq G$ with $|E(H)| \leq h$. For each such subgraph $H$, we denote $m(H)=|E(H)|$. We consider the following LP-relaxation of the (r,h)-Graph Partitioning problem, that has a variable $x_{H}$ for every graph $H \in \mathcal{H}$.

$$
\begin{array}{rcl}
\text { (LPW-P) } & & \\
\max & \sum_{H \in \mathcal{H}} m(H) \cdot x_{H} & \\
\text { s.t. } & & \\
& \sum_{\substack{H \in \mathcal{H}: \\
v \in V(H)}} x_{H} \leq 1 \quad \forall v \in V(G) \\
& \sum_{H \in \mathcal{H}} x_{H} \leq r & \\
& x_{H} \geq 0 & \forall H \in \mathcal{H}
\end{array}
$$

It is easy to verify that (LPW-P) is an LP-relaxation of the ( $r, h$ )-Graph Partitioning problem. Indeed, consider a solution $\left(H_{1}, \ldots, H_{r}\right)$ to the input instance $\operatorname{GP}(G, r, h)$. For all $1 \leq i \leq r$, we set $x_{H_{i}}=1$, and for every other graph $H \in \mathcal{H}$, we set $x_{H}=0$. This provides a feasible solution to (LPW-P),
whose value is precisely $\sum_{i=1}^{r}|E(H)|$. We denote the value of the optimal solution to (LPW-P) by $\mathrm{OPT}_{\mathrm{LP}-\mathrm{P}}$. From the above discussion, $\mathrm{OPT}_{\mathrm{LP}-\mathrm{P}} \geq \mathrm{OPT}_{\mathrm{GP}}(G, r, h)$.

In the following claim we provide an LP-rounding algorithm for (LPW-P). The claim is an analogue of Claim 4.2, Its proof is almost identical and is provided here for completeness.

Claim B.1. There is an efficient randomized algorithm, whose input consists of an instance $\operatorname{GP}(G, r, h)$ of the (r,h)-Graph Partitioning problem with $N=|V(G)|$, such that $N$ is greater than a large enough constant, and a solution $\left\{x_{H} \mid H \in \mathcal{H}\right\}$ to (LPW-P), in which the number of variables $x_{H}$ with $x_{H}>0$ is bounded by $O(\operatorname{poly}(N))$, and $\sum_{H \in \mathcal{H}} m(H) \cdot x_{H} \geq O P T_{\mathrm{LP}-\mathrm{P}} / \beta$, for some parameter $1 \leq \beta \leq N^{3}$; the solution is given by only specifying values of variables $x_{H}$ that are non-zero. The algorithm with high probability returns an integral solution $\left(H_{1}, \ldots, H_{r}\right)$ to instance $\mathrm{GP}(G, r, h)$, such that $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right| \geq \frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{2000 \beta \log ^{3} N}$.

Proof: We assume that we are given a solution $\left\{x_{H} \mid H \in \mathcal{H}\right\}$ to (LPW-P), in which the number of variables $x_{H}$ with $x_{H}>0$ is bounded by $O(\operatorname{poly}(N))$. Denote $C=\sum_{H \in \mathcal{H}} m(H) \cdot x_{H}$, and recall that $C \geq \mathrm{OPT}_{\mathrm{LP}-\mathrm{P}} / \beta$ holds. We denote by $\mathcal{H}^{\prime} \subseteq \mathcal{H}$ the collection of all graphs $H \in \mathcal{H}$ with $x_{H}>0$.

We construct another collection $\mathcal{H}^{\prime \prime} \subseteq \mathcal{H}^{\prime}$ of subgraphs of $G$ as follows. For every subgraph $H \in \mathcal{H}^{\prime}$, we add $H$ to $\mathcal{H}^{\prime \prime}$ independently, with probability $x_{H}$. Clearly, $\mathbf{E}\left[\sum_{H \in \mathcal{H}^{\prime \prime}} m(H)\right]=\sum_{H \in \mathcal{H}} m(H)$. $x_{H}=C$.

We say that a bad event $\mathcal{E}_{1}$ happens if some vertex $v \in V(G)$ lies in more than $5 \log N$ graphs in $\mathcal{H}^{\prime \prime}$. We say that a bad event $\mathcal{E}_{2}$ happens if $\left|\mathcal{H}^{\prime \prime}\right|>5 r \log N$. We say that a bad event $\mathcal{E}_{3}$ happens if $\sum_{H \in \mathcal{H}^{\prime \prime}} m(H)<\frac{C}{8}$. Lastly, we say that a bad event $\mathcal{E}$ happens if either of the events $\mathcal{E}_{1}, \mathcal{E}_{2}$, or $\mathcal{E}_{3}$ happen. The following observation is an analogue of Observation 4.3. Its proof is identical and is omitted here.

Observation B.2. $\operatorname{Pr}[\mathcal{E}] \leq 2 / N^{3}$.
Observe that we can efficiently check whether Event $\mathcal{E}$ happened. If Event $\mathcal{E}$ happens, then we terminate the algorithm with a FAIL. We assume from now on that Event $\mathcal{E}$ did not happen. Then $\sum_{H \in \mathcal{H}^{\prime \prime}} m(H) \geq \frac{C}{8} \geq \frac{\operatorname{OPT}_{G P}(G, r, h)}{8 \beta}$ must hold. We denote $\mathcal{H}^{\prime \prime}=\left\{H_{1}, H_{2}, \ldots, H_{z}\right\}$, where the graphs are indexed according to their value $m(H)$, so that $m\left(H_{1}\right) \geq m\left(H_{2}\right) \geq \cdots \geq m\left(H_{z}\right)$. We then let $\mathcal{H}^{*}=\left\{H_{1}, \ldots, H_{r}\right\}$ (if $z<r$, then we set $H_{z+1}=\cdots=H_{r}=\emptyset$ ). For all $1 \leq i \leq r$, we denote $E_{i}=E\left(H_{i}\right)$, so $\left|E_{i}\right|=m\left(H_{i}\right)$. Recall that, since Event $\mathcal{E}$ did not happen, $\left|\mathcal{H}^{\prime \prime}\right| \leq 5 r \log N$ holds. Therefore:

$$
\sum_{i=1}^{r}\left|E_{i}\right| \geq \frac{\sum_{H \in \mathcal{H}^{\prime \prime}} m(H)}{5 \log N} \geq \frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{40 \beta \log N}
$$

As before, the graphs in set $\mathcal{H}$ may not be mutually disjoint. However, since Event $\mathcal{E}$ did not happen, every vertex of $V(G)$ may lie in at most $5 \log N$ such graphs. We now construct a new collection $\mathcal{H}^{* *}=\left\{H_{1}^{\prime}, \ldots, H_{r}^{\prime}\right\}$ of graphs, as follows. For all $1 \leq i \leq r$, we will define a subset $V_{i} \subseteq V\left(H_{i}\right)$ of vertices, and we will then set $H_{i}^{\prime}=H_{i}\left[V_{i}\right]$. In order to define vertex sets $V_{1}, \ldots, V_{r}$, we start by setting $V_{1}=V_{2}=\cdots=V_{r}=\emptyset$, and then process vertices $v \in V(H)$ one by one. Consider any vertex $v \in V(G)$, and let $H_{i_{1}}, H_{i_{2}}, \ldots, H_{i_{a}} \in \mathcal{H}^{*}$ be the graphs of $\mathcal{H}^{*}$ containing $v$.

Vertex $v$ chooses an index $i^{*} \in\left\{i_{1}, \ldots, i_{a}\right\}$ at random, and is then added to $V_{i^{*}}$. Once all vertices of $V(G)$ are processed, we obtain a final collection $V_{1}, \ldots, V_{r}$ of sets of vertices, where for all $1 \leq i \leq r$, $V_{i} \subseteq V\left[H_{i}\right]$. For all $1 \leq i \leq r$, we then set $H_{i}^{\prime}=H_{i}\left[V_{i}\right]$. Since $H_{i}^{\prime} \subseteq H_{i}$, we are then guaranteed that $\left|E\left(H_{i}^{\prime}\right)\right| \leq h$ holds.
Note that for all $1 \leq j \leq r$, for every vertex $v \in V\left(H_{j}\right)$, the probability that $v \in V_{j}$ is at least $1 /(5 \log N)$. We say that an edge $e=(u, v) \in E_{j}$ survives if both $u, v \in V_{j}$. We denote by $E^{\prime \prime} \subseteq$ $\bigcup_{i=1}^{r} E_{i}$ the set of all edges that survive. Since $\operatorname{Pr}\left[u \in V_{j}\right] \geq 1 /(5 \log N), \operatorname{Pr}\left[v \in V_{j}\right] \geq 1 /(5 \log N)$, and the two events are independent, we get that the probability that edge $e$ survives is at least $1 /\left(25 \log ^{2} N\right)$. Overall, we get that:

$$
\mathbf{E}\left[\left|E^{\prime \prime}\right|\right] \geq \frac{\sum_{i=1}^{r}\left|E_{i}\right|}{25 \log ^{2} N} \geq \frac{\operatorname{OPT}_{\mathrm{GP}}(G, r, h)}{1000 \beta \log ^{3} N} .
$$

The final solution to instance $\operatorname{GP}(G, r, h)$ is $\mathcal{H}^{* *}=\left\{H_{1}^{\prime}, \ldots, H_{r}^{\prime}\right\}$. Clearly, the value of this solution is $\left|E^{\prime \prime}\right|$.

So far we have obtained a randomized algorithm that either returns FAIL (with probability at most $2 / N^{3}$ ), or it returns a solution to instance instance $\operatorname{GP}(G, r, h)$ of the ( $\mathrm{r}, \mathrm{h}$ )-Graph Partitioning problem, whose expected value is at least $\frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{1000 \beta \log ^{3} N}$.

Let $p^{\prime}$ be the probability that the algorithm returned a solution of value at least $\frac{\operatorname{OPT}_{G P}(G, r, h)}{2000 \beta \log ^{3} N}$, given that it did not return FAIL. Note that the expected solution value, assuming the algorithm did not return FAIL, is at most $\frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{2000 \beta \log ^{3} N}+p^{\prime} \cdot \mathrm{OPT}_{\mathrm{GP}}(G, r, h)$. Since this expectation is also at least $\frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{1000 \beta \log ^{3} N}$, we get that $p^{\prime} \geq \frac{1}{1000 \beta \log ^{3} N}$. Overall, the probability that our algorithm successfully returns a solution of value at least $\frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{2000 \beta \log ^{3} N}$ is $p^{\prime} \cdot \operatorname{Pr}[\neg \mathcal{E}] \geq \Omega\left(\frac{1}{\beta \log ^{3} N}\right)$. By repeating the algorithm $\operatorname{poly}(N)$ times we can ensure that it successfully computes a solution of value at least $\frac{\operatorname{OPT}_{\mathrm{GP}}(G, r, h)}{2000 \beta \log ^{3} N}$ with high probability.

## B. 2 Approximately Solving the LP-Relaxation

In this subsection we provide an approximate separation oracle for the dual linear program of (LPW-P). This is sufficient in order to obtain an algorithm with running time $O(\operatorname{poly}(N)$ ) that approximately solves (LPW-P) using the methods described in Section 4.2. The following Linear Program is a Dual of (LPW-P). It has a variable $y_{v}$ for every vertex $v \in V(G)$, and an additional variable $z$.

$$
\begin{array}{ccc}
\text { (LPW-D) } & \\
\text { min } & r \cdot z+\sum_{v \in V(G)} y_{v} & \\
\text { s.t. } & \\
& z+\sum_{v \in V(H)} y_{v} \geq m(H) & \forall H \in \mathcal{H} \\
& z \geq 0 & \\
& y_{v} \geq 0 & \forall v \in V(G)
\end{array}
$$

We denote the value of the optimal solution to (LPW-D) by $\mathrm{OPT}_{\mathrm{LPW}}$.D.
The following lemma provides a randomized separation oracle for (LPW-D). It is an analogue of Lemma 4.4, and its proof is essentially identical. We provide it here for completeness.

Lemma B.3. Assume that there is an efficient $\alpha(n)$-approximation algorithm for the Densest $k$ Subgraph problem, where $\alpha$ is an increasing function, and $n$ is the number of vertices in the input graph. Then there is a randomized $\beta(N)$-approximate separation oracle for ( $L P W-D$ ), where $N$ is the number of variables in the input graph $G$, and $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$.

Proof: Recall that we are given as input real values $z$ and $\left\{y_{v} \mid v \in V(G)\right\}$. As before, we can efficiently check whether $z \geq 0$, and whether $y_{v} \geq 0$ for all $v \in V(G)$. If this is not the case, we can return the corresponding violated constraint.

We say that a subgraph $H \in \mathcal{H}$ is bad if $z+\sum_{v \in V(H)} y_{v}<m(H) / \beta$ holds, where $\beta=c \cdot \alpha\left(N^{2}\right) \cdot \log ^{2} N$, and $c$ is a large enough constant whose value we set later. Our goal is to design an efficient algorithm that either returns a violated constraint of the LP (that is, a graph $H \in \mathcal{H}$ for which $z+\sum_{v \in V(H)} y_{v}<m(H)$ holds); or it returns "accept". We require that, if there exists a bad subgraph $H \in \mathcal{H}$, then the probability that the algorithm returns "accept" is at most $2 / 3$.
We slightly modify the input values in $\left\{y_{v} \mid v \in V(G)\right\}$, almost exactly like in the proof of Lemma B.3. First, for every vertex $v \in V(G)$ with $y_{v}>h$, we let $y_{v}^{\prime}$ be the smallest integral power of 2 that is greater than $h$, and for every vertex $v \in V(G)$ with $y_{v}<1 / 4$, we set $y_{v}^{\prime}=0$. For each remaining vertex $v$, we let $y_{v}^{\prime}$ be the smallest integral power of 2 that is greater than $4 y_{v}$. Notice that, for every vertex $v$ with $y_{v}^{\prime} \neq 0,1 \leq y_{v}^{\prime} \leq 4 h$ holds, and $y_{v}^{\prime}$ is an integral power of 2 . We also set $z^{\prime}=2 z$. We say that a subgraph $H \in \mathcal{H}$ is problematic if $z^{\prime}+\sum_{v \in V(H)} y_{v}^{\prime}<8 m(H) / \beta$ holds. We use the following two observations, that are analogues of Observation 4.5 and Observation 4.6 their proofs are also almost identical.

Observation B.4. If $H \in \mathcal{H}$ is a bad subgraph of $G$, then it is a problematic subgraph of $G$.

Proof: Recall that, if $H$ is a bad subgraph, then $z+\sum_{v \in V(H)} y_{v}<m(H) / \beta$ must hold. Since, for every vertex $v \in V(G), y_{v}^{\prime} \leq 8 y_{v}$, and $z^{\prime}=2 z$, we get that:

$$
z^{\prime}+\sum_{v \in V(H)} y_{v}^{\prime} \leq 2 z+8 \sum_{v \in V(H)} y_{v} \leq 8\left(z+\sum_{v \in V(H)} y_{v}\right)<8 m(H) / \beta .
$$

Therefore, subgraph $H$ is problematic.
Observation B.5. Assume that there exists a subgraph $H \in \mathcal{H}$, for which $z^{\prime}+\sum_{v \in V(H)} y_{v}^{\prime}<m(H)$ holds. Let $H^{\prime} \subseteq H$ be the graph obtained from $H$ after removing all isolated vertices from it. Then $z+\sum_{v \in V\left(H^{\prime}\right)} y_{v}<m\left(H^{\prime}\right)$ holds.

Proof: Since every vertex $v \in V(H) \backslash V\left(H^{\prime}\right)$ is isolated in $H$, we get that $m\left(H^{\prime}\right)=\left|E\left(H^{\prime}\right)\right|=$ $|E(H)|=m(H)$. We partition the vertices of $H^{\prime}$ into two subsets: set $X$ containing all vertices $v \in V\left(H^{\prime}\right)$ with $y_{v}<1 / 4$, and set $Y$ containing all remaining vertices. Clearly, $\sum_{v \in X} y_{v}<\frac{|X|}{4} \leq$ $\frac{m\left(H^{\prime}\right)}{2}$ (since $m\left(H^{\prime}\right) \geq\left|V\left(H^{\prime}\right)\right| / 2 \geq|X| / 2$, as graph $H^{\prime}$ contains no isolated vertices).

Assume for contradiction that $z+\sum_{v \in V\left(H^{\prime}\right)} y_{v} \geq m\left(H^{\prime}\right)$. Then:

$$
z+\sum_{v \in Y} y_{v} \geq m\left(H^{\prime}\right)-\sum_{v \in X} y_{v} \geq m\left(H^{\prime}\right) / 2 .
$$

We now consider two cases. The first case is when there is some vertex $v \in Y$ with $y_{v} \geq h$. In this case, $y_{v}^{\prime} \geq h$ holds, and $z^{\prime}+\sum_{v \in S} y_{v}^{\prime} \geq h>m\left(H^{\prime}\right)$ holds, a contradiction.
Otherwise, for every vertex $v \in Y, y_{v}^{\prime} \geq 4 y_{v}$ holds. Since $z^{\prime}=2 z$ also holds, we get that:

$$
z^{\prime}+\sum_{v \in V(H)} y_{v}^{\prime} \geq z^{\prime}+\sum_{v \in Y} y_{v}^{\prime} \geq 2 z+4 \sum_{v \in Y} y_{v} \geq m\left(H^{\prime}\right)=m(H),
$$

a contradiction.
From now on we focus on values $z^{\prime},\left\{y_{v}^{\prime} \mid v \in V(G)\right\}$. It is now enough to design an efficient randomized algorithm, that either computes a subgraph $H \in \mathcal{H}$, for which $z^{\prime}+\sum_{v \in V(H)} y_{v}^{\prime}<m(H)$ holds, or returns "accept". It is enough to ensure that, if there is a problematic subraph $H \in \mathcal{H}$, then the algorithm returns "accept" with probability at most $2 / 3$. Indeed, if there is a bad subgraph $H \in \mathcal{H}$, then, from Observation B.4, there is a problematic subgraph, and the algorithm will return "accept" with probability at most $2 / 3$. On the other hand, if the algorithm computes a subgraph $H \in \mathcal{H}$ of vertices, for which $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<m(S)$ holds, then we can return the subgraph $H^{\prime} \subseteq H$ from the statement of Observation B.5, that defines a violated constraint with respect to the original LP-values.

Our algorithm is essentially the same as before: it computes a random partition $(A, B)$ of the vertices of $G$, where every vertex $v \in V(G)$ is independently added to $A$ or to $B$ with probability $1 / 2$ each. Let $q=\lceil\log (8 h)\rceil$. For all $1 \leq i \leq q$, we define a set $A_{i} \subseteq A$ of vertices: $A_{i}=\{v \in A \mid$ $\left.y_{v}^{\prime}=2^{i-1}\right\}$, and we let $A_{0}=\left\{v \in A \mid y_{v}^{\prime}=0\right\}$. Clearly, $\left(A_{0}, \ldots, A_{q}\right)$ is a partition of the set $A$ of vertices.

We compute a partition $\left(B_{0}, \ldots, B_{q}\right)$ of the vertices of $B$ similarly. For all $0 \leq i, j \leq q$, we denote by $E_{i, j}$ the set of all edges $e=(u, v)$ with $u \in A_{i}$ and $v \in B_{j}$, and we define a bipartite graph $G_{i, j}$, whose vertex set is $A_{i} \cup B_{j}$, and edge set is $E_{i, j}$.

Recall that we have assumed that there is an efficient $\alpha(n)$-approximation algorithm for the Densest $k$-Subgraph problem, where $n$ is the number of vertices in the input graph. From Lemma 2.1, there exists an efficient $O\left(\alpha\left(\hat{n}^{2}\right)\right)$-approximation algorithm for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem, where $\hat{n}$ is the number of vertices in the input graph. We denote this algorithm by $\mathcal{A}^{\prime}$.
For every pair $0 \leq i, j \leq q$ of integers, and every pair $0 \leq k_{1}, k_{2} \leq N$ of integers, we apply Algorithm $\mathcal{A}^{\prime}$ for the Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph problem to graph $G_{i, j}$, with parameters $k_{1}$ and $k_{2}$. Let $S_{i, j}^{k_{1}, k_{2}}$ be the output of this algorithm, and let $m_{i, j}^{k_{1}, k_{2}}$ be the number of edges in the subgraph of $G_{i, j}$ that is induced by the set $S_{i, j}^{k_{1}, k_{2}}$ of vertices. We say that the application of algorithm $\mathcal{A}^{\prime}$ is successful if $z^{\prime}+\sum_{v \in S_{i, j}^{k_{1}, k_{2}}} y_{v}^{\prime}<\min \left\{h, m_{i, j}^{k_{1}, k_{2}}\right\}$, and otherwise it is unsuccessful. If, for any quadruple ( $i, j, k_{1}, k_{2}$ ) of indices, the application of algorithm $\mathcal{A}^{\prime}$ was successful, then we return a graph $H$, that is defined as the subgraph of $G$ induced by the set $S=S_{i, j}^{k_{1}, k_{2}}$ of vertices; if this graph contains more than $h$ edges, then we delete arbitrary edges from it, until $|E(H)|=h$ holds.

Clearly, $H \in \mathcal{H}$ must hold. Moreover, we are guaranteed that $z^{\prime}+\sum_{v \in V(H)} y_{v}^{\prime}<\min \left\{h, m_{i, j}^{k_{1}, k_{2}}\right\} \leq$ $m(H)$, as required. If every application of algorithm $\mathcal{A}^{\prime}$ is unsuccessful, then we return "accept". The following observation will finish the proof of Lemma B.3. The observation is an analogue of Observation 4.7 and its proof is essentially identical.

Observation B.6. Suppose there is a problematic subgraph $H \in \mathcal{H}$. Then the probability that the algorithm returns "accept" is at most $2 / 3$.

Proof: Let $H \in \mathcal{H}$ be a problematic subgraph, and denote $S=V(H)$, so $z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<8 m(S) / \beta$ holds. Let $E^{\prime}=E(H)$, so $\left|E^{\prime}\right|=m(H) \leq h$.

Denote $A_{S}=A \cap S, B_{S}=B \cap S$, and let $E^{\prime \prime} \subseteq E^{\prime}$ be the set of edges $e$, such that exactly one endpoint of $e$ lies in $A$. Clearly, for every edge $e \in E^{\prime}, \operatorname{Pr}\left[e \in E^{\prime \prime}\right]=1 / 2$. Therefore, $\mathbf{E}\left[\left|E^{\prime \prime}\right|\right]=\left|E^{\prime}\right| / 2$. Let $\mathcal{E}^{\prime}$ be the bad event that $\left|E^{\prime \prime}\right|<\left|E^{\prime}\right| / 8$. Using the same arguments as in the proof of Observation 4.7, $\operatorname{Pr}\left[\mathcal{E}^{\prime}\right] \leq 2 / 3$. Next, we show that, if Event $\mathcal{E}^{\prime}$ does not happen, then the algorithm does not return "accept".
From now on we assume that Event $\mathcal{E}^{\prime}$ did not happen, so $\left|E^{\prime \prime}\right| \geq\left|E^{\prime}\right| / 8$. Therefore:

$$
z^{\prime}+\sum_{v \in S} y_{v}^{\prime}<\frac{8 m(H)}{\beta} \leq \frac{64\left|E^{\prime \prime}\right|}{\beta}
$$

holds.
Clearly, there must be a pair $0 \leq i, j \leq q$ of indices, such that $\left|E^{\prime \prime} \cap E_{i, j}\right| \geq \frac{\left|E^{\prime \prime}\right|}{4 q^{2}} \geq \frac{\left|E^{\prime \prime \prime}\right|}{128 \log ^{2} m}$. We now fix this pair $i, j$ of indices, and denote $A_{i}^{\prime}=A_{i} \cap S$ and $B_{j}^{\prime}=B_{j} \cap S$. We also denote $k_{1}=\left|A_{i}^{\prime}\right|$ and let $k_{2}=\left|B_{j}^{\prime}\right|$. Denote $M_{i, j}=\left|E^{\prime \prime} \cap E_{i, j}\right|$. Fom our choice of indices $i, j$, we get that:

$$
z^{\prime}+\sum_{v \in A_{i}^{\prime} \cup B_{j}^{\prime}} y_{v}^{\prime} \leq z^{\prime}+\sum_{v \in S} y_{v}^{\prime} \leq \frac{64\left|E^{\prime \prime}\right|}{\beta} \leq \frac{M_{i, j}}{\beta} \cdot\left(2^{13} \cdot \log ^{2} m\right) .
$$

Notice that the set $S^{\prime}=A_{i}^{\prime} \cup B_{j}^{\prime}$ of vertices provides a solution to the instance of the Bipartite Densest ( $k_{1}, k_{2}$ )-Subgraph problem on graph $G_{i, j}$ with parameters $k_{1}, k_{2}$, whose value is at least $M_{i, j}$. Let $S^{\prime \prime}=S_{i, j}^{k_{1}, k_{2}}$ be the set of vertices obtained by applying Algorithm $\mathcal{A}^{\prime}$ to graph $G_{i, j}$ with parameters $k_{1}, k_{2}$. Since $\left|V\left(G_{i, j}\right)\right| \leq N$, and since $\mathcal{A}^{\prime}$ is an $O\left(\alpha\left(N^{2}\right)\right)$-approximation algorithm for Bipartite Densest $\left(k_{1}, k_{2}\right)$-Subgraph, we are guaranteed that $\left|E_{G}\left(S^{\prime \prime}\right)\right| \geq \Omega\left(\frac{M_{i, j}}{\alpha\left(N^{2}\right)}\right)$. Recall that $\left|A \cap S^{\prime \prime}\right| \leq k_{1}$; $A \cap S^{\prime \prime} \subseteq A_{i}$, and all vertices $v \in A_{i}$ have an identical value $y_{v}^{\prime}$. Therefore, $\sum_{v \in A \cap S^{\prime \prime}} y_{v}^{\prime} \leq \sum_{v \in A_{i}^{\prime}} y_{v}^{\prime}$. Using a similar reasoning, $\sum_{v \in B \cap S^{\prime \prime}} y_{v}^{\prime} \leq \sum_{v \in B_{j}^{\prime}} y_{v}^{\prime}$. Overall, we then get that:

$$
\begin{aligned}
z^{\prime}+\sum_{v \in S^{\prime \prime}} y_{v}^{\prime} & \leq z^{\prime}+\sum_{v \in S^{\prime}} y_{v}^{\prime} \\
& \leq \frac{M_{i, j}}{\beta} \cdot\left(2^{13} \cdot \log ^{2} m\right) \\
& \leq O\left(\frac{\alpha\left(N^{2}\right) \cdot 2^{13} \cdot \log ^{2} m}{\beta}\right) \cdot \min \left\{\left|E_{G}\left(S^{\prime \prime}\right)\right|, h\right\} .
\end{aligned}
$$

Recall that $\beta=c \cdot \alpha\left(N^{2}\right) \cdot \log ^{2} N$. By letting the value of the constant $c$ be large enough, we can ensure that $z^{\prime}+\sum_{v \in S^{\prime \prime}} y_{v}^{\prime}<\min \left\{\left|E_{G}\left(S^{\prime \prime}\right)\right|, h\right\}$, and so the application of algorithm $\mathcal{A}^{\prime}$ to graph $G_{i, j}$ with parameters $k_{1}$ and $k_{2}$ is guaranteed to be successful. Therefore, if Event $\mathcal{E}^{\prime}$ does not happen, and we set $c$ to be a large enough constant, then our algorithm does not return "accept". Since $\operatorname{Pr}\left[\mathcal{E}^{\prime}\right] \leq 2 / 3$, the observation follows.
We can use the separation oracle described in Lemma B.3 in order to obtain a $\beta(N)$-approximate solution to (LPW-P), whose support size is bounded by $O(\operatorname{poly}(N))$ using the standard techniques that were described in Section 4.2; we do not repeat them here. By applying the LP-rounding algorithm from Claim B. 1 to the resulting LP-solution, with high probability we obtain, in time $O(\operatorname{poly}(N))$, an integral solution $\left(H_{1}, \ldots, H_{r}\right)$ to instance $\operatorname{GP}(G, r, h)$, such that $\sum_{i=1}^{r}\left|E\left(H_{i}\right)\right| \geq$ $\Omega\left(\frac{\mathrm{OPT}_{\mathrm{GP}}(G, r, h)}{300 \beta(N) \log ^{3} N}\right)$. Since $\beta(N)=O\left(\alpha\left(N^{2}\right) \cdot \log ^{2} N\right)$, with high probability we obtain an $O\left(\alpha\left(N^{2}\right)\right.$. poly $\log N$ )-approximate solution to the input instance of (r,h)-Graph Partitioning.

## References

$\left[\mathrm{AAM}^{+} 11\right]$ Noga Alon, Sanjeev Arora, Rajsekar Manokaran, Dana Moshkovitz, and Omri Weinstein. Inapproximabilty of densest k-subgraph from average case hardness. Manuscript, 2011. https://www.tau.ac.il/~nogaa/PDFS/dks8.pdf.
[ABS15] Sanjeev Arora, Boaz Barak, and David Steurer. Subexponential algorithms for unique games and related problems. Journal of the ACM (JACM), 62(5):1-25, 2015.
[ACNS82] M. Ajtai, V. Chvátal, M. Newborn, and E. Szemerédi. Crossing-free subgraphs. Theory and Practice of Combinatorics, pages 9-12, 1982.
[ADD $\left.{ }^{+} 93\right]$ Ingo Althöfer, Gautam Das, David Dobkin, Deborah Joseph, and José Soares. On sparse spanners of weighted graphs. Discrete $\mathcal{E}$ Computational Geometry, 9(1):81100, 1993.
[AMS07] Christoph Ambuhl, Monaldo Mastrolilli, and Ola Svensson. Inapproximability results for sparsest cut, optimal linear arrangement, and precedence constrained scheduling. In 48 th Annual IEEE Symposium on Foundations of Computer Science (FOCS'07), pages 329-337. IEEE, 2007.
[ARV09] Sanjeev Arora, Satish Rao, and Umesh V. Vazirani. Expander flows, geometric embeddings and graph partitioning. J. ACM, 56(2), 2009.
[Bar15] Siddharth Barman. Approximating nash equilibria and dense bipartite subgraphs via an approximate version of caratheodory's theorem. In Proceedings of the forty-seventh annual ACM symposium on Theory of computing, pages 361-369, 2015.
$\left[\mathrm{BCC}^{+} 10\right]$ Aditya Bhaskara, Moses Charikar, Eden Chlamtac, Uriel Feige, and Aravindan Vijayaraghavan. Detecting high log-densities: an $O\left(n^{1 / 4}\right)$ approximation for densest $k$-subgraph. In Proceedings of the $42 n d$ ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010, pages 201-210, 2010.
[ $\left.\mathrm{BCG}^{+} 12\right]$ Aditya Bhaskara, Moses Charikar, Venkatesan Guruswami, Aravindan Vijayaraghavan, and Yuan Zhou. Polynomial integrality gaps for strong sdp relaxations of densest k-subgraph. In Proceedings of the twenty-third annual ACM-SIAM symposium on Discrete Algorithms, pages 388-405. SIAM, 2012.
[Ber45] Joseph Bertrand. Memory on the number of values a function can take: when you swap the letters it contains. Bachelor, 1845.
[ $\left.\mathrm{BGH}^{+} 15\right]$ Boaz Barak, Parikshit Gopalan, Johan Håstad, Raghu Meka, Prasad Raghavendra, and David Steurer. Making the long code shorter. SIAM Journal on Computing, 44(5):1287-1324, 2015.
[BKRW17] Mark Braverman, Young Kun Ko, Aviad Rubinstein, and Omri Weinstein. Eth hardness for densest-k-subgraph with perfect completeness. In Proceedings of the TwentyEighth Annual ACM-SIAM Symposium on Discrete Algorithms, pages 1326-1341. SIAM, 2017.
[BKS19] Boaz Barak, Pravesh K. Kothari, and David Steurer. Small-set expansion in shortcode graph and the 2-to-2 conjecture. In 10th Innovations in Theoretical Computer Science Conference, ITCS 2019, January 10-12, 2019, San Diego, California, USA, pages 9:1-9:12, 2019.
[Bol04] Béla Bollobás. Extremal graph theory. Courier Corporation, 2004.
[Cab13] Sergio Cabello. Hardness of approximation for crossing number. Discrete E Computational Geometry, 49(2):348-358, 2013.
$\left[\mathrm{CCH}^{+} 20\right]$ Shih-Chia Chang, Li-Hsuan Chen, Ling-Ju Hung, Shih-Shun Kao, and Ralf Klasing. The hardness and approximation of the densest k-subgraph problem in parameterized metric graphs. In 2020 International Computer Symposium (ICS), pages 126-130. IEEE, 2020.
$\left[\mathrm{CDK}^{+} 18\right]$ Eden Chlamtác, Michael Dinitz, Christian Konrad, Guy Kortsarz, and George Rabanca. The densest k-subhypergraph problem. SIAM Journal on Discrete Mathematics, 32(2):1458-1477, 2018.
[Čeb50] Pafnutij Lvovič Čebyšev. Mémoire sur les nombres premiers. 1850.
[CH11] Markus Chimani and Petr Hliněnỳ. A tighter insertion-based approximation of the crossing number. In International Colloquium on Automata, Languages, and Programming, pages 122-134. Springer, 2011.
[Chu11] Julia Chuzhoy. An algorithm for the graph crossing number problem. In Proceedings of the forty-third annual ACM symposium on Theory of computing, pages 303-312. ACM, 2011.
[Chu15] Julia Chuzhoy. Excluded grid theorem: Improved and simplified. In Proceedings of the forty-seventh annual ACM symposium on Theory of Computing, pages 645-654, 2015.
[CKN21] Julia Chuzhoy, David Hong Kyun Kim, and Rachit Nimavat. Almost polynomial hardness of node-disjoint paths in grids. Theory of Computing, 17(6):1-57, 2021.
[CMS11] Julia Chuzhoy, Yury Makarychev, and Anastasios Sidiropoulos. On graph crossing number and edge planarization. In Proceedings of the twenty-second annual ACMSIAM symposium on Discrete algorithms, pages 1050-1069. SIAM, 2011.
[CMT20] Julia Chuzhoy, Sepideh Mahabadi, and Zihan Tan. Towards better approximation of graph crossing number. In 2020 IEEE 61st Annual Symposium on Foundations of Computer Science (FOCS), pages 73-84. IEEE, 2020. Full version: Arxiv:2011.06545.
[CS13] Chandra Chekuri and Anastasios Sidiropoulos. Approximation algorithms for euler genus and related problems. In 2013 IEEE 54th Annual Symposium on Foundations of Computer Science, pages 167-176. IEEE, 2013.
[CT22] Julia Chuzhoy and Zihan Tan. A subpolynomial approximation algorithm for graph crossing number in low-degree graphs. In Proceedings of the 54th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2022, pages 303-316, 2022. Full version: Arxiv:2202.06827.
[DKK ${ }^{+}$18a] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. On non-optimally expanding sets in Grassmann graphs. In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 25-29, 2018, pages 940-951, 2018.
[DKK ${ }^{+}$18b] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. Towards a proof of the 2-to-1 games conjecture? In Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, Los Angeles, CA, USA, June 2529, 2018, pages 376-389, 2018.
[DP09] Devdatt P Dubhashi and Alessandro Panconesi. Concentration of measure for the analysis of randomized algorithms. Cambridge University Press, 2009.
[EGS02] Guy Even, Sudipto Guha, and Baruch Schieber. Improved approximations of crossings in graph drawings and vlsi layout areas. SIAM Journal on Computing, 32(1):231-252, 2002.
[Fei02] Uriel Feige. Relations between average case complexity and approximation complexity. In Proceedings of the thiry-fourth annual ACM symposium on Theory of computing, pages 534-543, 2002.
[FL01] Uriel Feige and Michael Langberg. Approximation algorithms for maximization problems arising in graph partitioning. Journal of Algorithms, 41(2):174-211, 2001.
[FPK01] Uriel Feige, David Peleg, and Guy Kortsarz. The dense k-subgraph problem. Algorithmica, 29(3):410-421, 2001.
[FS $\left.{ }^{+} 97\right]$ Uriel Feige, Michael Seltser, et al. On the densest k-subgraph problem. Technical Report CS97-16, Weizmann Institute of Science., 1997. https://citeseerx.ist. psu.edu/viewdoc/download?doi=10.1.1.37.9962\&rep=rep1\&type=pdf.
[GL09] Doron Goldstein and Michael Langberg. The dense k subgraph problem. arXiv preprint arXiv:0912.5327, 2009.
[Han22] Tesshu Hanaka. Computing densest $k$-subgraph with structural parameters. arXiv preprint arXiv:2207.09803, 2022.
[Hås01] Johan Håstad. Some optimal inapproximability results. Journal of the ACM (JACM), 48(4):798-859, 2001.
[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In Proceedings of the thiry-fourth annual ACM symposium on Theory of computing, pages 767-775, 2002.
[Kho06] Subhash Khot. Ruling out ptas for graph min-bisection, dense k-subgraph, and bipartite clique. SIAM Journal on Computing, 36(4):1025-1071, 2006.
[KLS00] Sanjeev Khanna, Nathan Linial, and Shmuel Safra. On the hardness of approximating the chromatic number. Combinatorica, 20(3):393-415, 2000.
[KMMS18] Subhash Khot, Dor Minzer, Dana Moshkovitz, and Muli Safra. Small set expansion in the Johnson graph. Electronic Colloquium on Computational Complexity (ECCC), 25:78, 2018. https://eccc.weizmann.ac.il/report/2018/078.
[KMS17] Subhash Khot, Dor Minzer, and Muli Safra. On independent sets, 2-to-2 games, and Grassmann graphs. In Proceedings of the 49 th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pages 576-589, 2017.
[KMS18] Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in Grassmann graph have near-perfect expansion. In 59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018, Paris, France, October 7-9, 2018, pages 592-601, 2018.
[KP93] G Kortsarz and D Peleg. On choosing a dense subgraph. In Proceedings of 1993 IEEE 34th Annual Foundations of Computer Science, pages 692-701. IEEE Computer Society, 1993.
[KS13] Subhash Khot and Muli Safra. A two-prover one-round game with strong soundness. Theory of Computing, 9:863-887, 2013.
[KS17] Ken-ichi Kawarabayashi and Anastasios Sidiropoulos. Polylogarithmic approximation for minimum planarization (almost). In 58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017, Berkeley, CA, USA, October 15-17, 2017, pages 779-788, 2017.
[KS19] Ken-ichi Kawarabayashi and Anastasios Sidiropoulos. Polylogarithmic approximation for euler genus on bounded degree graphs. In Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, pages 164-175. ACM, 2019.
[Lei83] F. T. Leighton. Complexity issues in VLSI: optimal layouts for the shuffle-exchange graph and other networks. MIT Press, 1983.
[Lin18] Bingkai Lin. The parameterized complexity of the k-biclique problem. Journal of the ACM (JACM), 65(5):1-23, 2018.
[LR99] Tom Leighton and Satish Rao. Multicommodity max-flow min-cut theorems and their use in designing approximation algorithms. Journal of the ACM (JACM), 46(6):787832, 1999.
[LT79] Richard J Lipton and Robert Endre Tarjan. A separator theorem for planar graphs. SIAM Journal on Applied Mathematics, 36(2):177-189, 1979.
[Man17] Pasin Manurangsi. Almost-polynomial ratio ETH-hardness of approximating densest k-subgraph. In Proceedings of the 49 th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, Montreal, QC, Canada, June 19-23, 2017, pages 954-961, 2017.
[Man18] Pasin Manurangsi. Inapproximability of maximum biclique problems, minimum k-cut and densest at-least-k-subgraph from the small set expansion hypothesis. Algorithms, 11(1):10, 2018.
[Mat02] J. Matoušek. Lectures on discrete geometry. Springer-Verlag, 2002.
[PSS96] János Pach, Farhad Shahrokhi, and Mario Szegedy. Applications of the crossing number. Algorithmica, 16(1):111-117, 1996.
[PT00] J. Pach and G. Tóth. Thirteen problems on crossing numbers. Geombinatorics, 9(4):194-207, 2000.
[RS09] R. B. Richter and G. Salazar. Crossing numbers. In L. W. Beineke and R. J. Wilson, editors, Topics in Topological Graph Theory, chapter 7, pages 133-150. Cambridge University Press, 2009.
[Sch12] Marcus Schaefer. The graph crossing number and its variants: A survey. The electronic journal of combinatorics, pages DS21-Sep, 2012.
[Sot20] Renata Sotirov. On solving the densest k-subgraph problem on large graphs. Optimization Methods and Software, 35(6):1160-1178, 2020.
[Ste] David Steurer. Subexponential algorithms for d-to-1 two-prover games and for certifying almost perfect expansion. Available at https://citeseerx.ist.psu.edu/ viewdoc/download?doi=10.1.1.189.5388\&rep=rep1\&type=pdf, 2010.
[Tur77] P. Turán. A note of welcome. J. Graph Theory, 1:1-5, 1977.


[^0]:    *Toyota Technological Institute at Chicago. Email: cjulia@ttic.edu. Supported in part by NSF grant CCF2006464.
    ${ }^{\dagger}$ Massachusetts Institute of Technology. Email: minad@mit.edu. Part of the work was done while the author was a summer intern at TTIC.
    ${ }^{\ddagger}$ Weizmann Institute of Science. Email: vadim.grinberg@weizmann.ac.il.
    ${ }^{\S}$ DIMACS, Rutgers University. Email: zihantan1993@gmail.com. Supported by a grant to DIMACS from the Simons Foundation (820931). Work done while the author was a graduate student at University of Chicago.

