

Almost Optimal Inapproximability of Multidimensional Packing Problems

Sai Sandeep*

Computer Science Department
Carnegie Mellon University
Pittsburgh, PA 15213

Abstract

Multidimensional packing problems generalize the classical packing problems such as Bin Packing, Multiprocessor Scheduling by allowing the jobs to be d -dimensional vectors. While the approximability of the scalar problems is well understood, there has been a significant gap between the approximation algorithms and the hardness results for the multidimensional variants. In this paper, we close this gap by giving almost tight hardness results for these problems.

1. We show that Vector Bin Packing has no $\Omega(\log d)$ factor asymptotic approximation algorithm when d is a large constant, assuming $P \neq NP$. This matches the $\ln d + O(1)$ factor approximation algorithms (Chekuri, Khanna SICOMP 2004, Bansal, Caprara, Sviridenko SICOMP 2009, Bansal, Eliáš, Khan SODA 2016) upto constants.
2. We show that Vector Scheduling has no polynomial time algorithm with an approximation ratio of $\Omega((\log d)^{1-\epsilon})$ when d is part of the input, assuming $NP \not\subseteq ZPTIME(n^{(\log n)^{O(1)}})$. This almost matches the $O\left(\frac{\log d}{\log \log d}\right)$ factor algorithms (Harris, Srinivasan JACM 2019, Im, Kell, Kulkarni, Panigrahi SICOMP 2019). We also show that the problem is NP-hard to approximate within $(\log \log d)^{\omega(1)}$.
3. We show that Vector Bin Covering is NP-hard to approximate within $\Omega\left(\frac{\log d}{\log \log d}\right)$ when d is part of the input, almost matching the $O(\log d)$ factor algorithm (Alon *et al.*, Algorithmica 1998).

Previously, no hardness results that grow with d were known for Vector Scheduling and Vector Bin Covering when d is part of the input and for Vector Bin Packing when d is a fixed constant.

*spaller1@andrew.cmu.edu. Research supported in part by NSF grants CCF-1563742 and CCF-1908125.

1 Introduction

Bin Packing and Multiprocessor Scheduling (also known as Makespan Minimization) are one of the most fundamental problems in Combinatorial Optimization. They have been studied intensely from the early days of approximation algorithms and have had a great impact on the field. These two are packing problems where we have n jobs with certain sizes, and the objective is to pack them into bins efficiently. In Bin Packing, each bin has unit size and the objective is to minimize the number of bins, while in Multiprocessor Scheduling, we are given a fixed number of bins and the objective is to minimize the maximum load in a bin. A relatively less studied but natural problem is Bin Covering, where the objective is to assign the jobs to the maximum number of bins such that in each bin, the load is at least 1. These problems are well understood in terms of approximation algorithms: all the three problems are NP-hard, and all of them have a Polynomial Time Approximation Scheme (PTAS) [dVLL81, HS87, CJK01].

In this paper, we study the approximability of the multidimensional generalizations of these three problems. The three corresponding problems are Vector Bin Packing, Vector Scheduling, and Vector Bin Covering. Apart from their theoretical importance, these problems are widely applicable in practice [Spi94, ST12, PTUW11] where the jobs often have multiple dimensions such as CPU, Hard disk, memory, etc.

In the Vector Bin Packing problem, the input is a set of n vectors in $[0, 1]^d$ and the goal is to partition the vectors into the minimum number of parts such that in each part, the sum of vectors is at most 1 in every coordinate. The problem behaves differently from Bin Packing even when $d = 2$: Woeginger [Woe97] proved that there is no asymptotic¹ PTAS for 2-dimensional Vector Bin Packing, assuming $P \neq NP$. On the algorithmic front, the PTAS for Bin Packing [dVLL81] easily implies a $d + \epsilon$ approximation for Vector Bin Packing. When d is part of the input, this is almost tight: there is a lower bound of $d^{1-\epsilon}$ shown by [CK04]². When d is a fixed constant³, much better algorithms are known [CK04, BCS09, BEK16] that get $\ln d + O(1)$ approximation guarantee. However, the best hardness factor (for arbitrary constant d) is still the APX-hardness result of the 2-dimensional problem due to Woeginger from 1997. Closing this gap, either by obtaining a $O(1)$ factor algorithm or showing a hardness factor that is a function of d , has remained a challenging open problem. It is one of the ten open problems in a recent survey on multidimensional scheduling problems [CKPT17]. It also appeared in a recent report by Bansal [Ban17] on open problems in scheduling.

In the Vector Scheduling problem, given a set of n vector jobs in $[0, 1]^d$, and m identical machines, the objective is to assign the jobs to machines to minimize the maximum ℓ_∞ norm of the load on the machines. Chekuri and Khanna [CK04] introduced the problem as a natural generalization of Multiprocessor Scheduling and obtained a PTAS for the problem when d is a fixed constant. When d is part of the input, they obtained a $O(\log^2 d)$ factor approximation algorithm. They also showed that it is NP-hard to obtain a C factor approximation algorithm for the problem, for any constant C . Meyerson, Roytman, and Tagiku [MRT13] gave an improved $O(\log d)$ factor algorithm while the current best factor is $O\left(\frac{\log d}{\log \log d}\right)$ due to Harris and Srinivasan [HS19] and Im, Kell, Kulkarni, and Panigrahi [IKKP19]. The algorithm of Harris and Srinivasan [HS19] works for the more general setting of unrelated machines where each job can have a different vector load for each machine. However, no super constant hardness is known even in this unrelated

¹The asymptotic approximation ratio (formally defined in Section 2) of an algorithm is the ratio of its cost and the optimal cost when the optimal cost is large enough. All the approximation factors mentioned in this paper for Vector Bin Packing are asymptotic.

²[CK04] actually give $d^{\frac{1}{2}-\epsilon}$ hardness, but it has been shown later (see e.g., [CKPT17]) that a slight modification of their reduction gives $d^{1-\epsilon}$ hardness.

³The algorithms are now allowed to run in time $n^{f(d)}$, for some function f .

machines setting.

In the Vector Bin Covering problem, the input is a set of n vectors in $[0, 1]^d$. The objective is to partition these into the maximum number of parts such that in each part, the sum of vectors is at least 1 in every coordinate. It is also referred to as “dual Vector Packing” in the literature. This problem is introduced by Alon *et al.* [AAC⁺98] who gave a $O(\log d)$ factor approximation algorithm. They also gave a d factor algorithm using a method from the area of compact vector approximation that outperforms the above algorithm for small values of d . On the hardness front, Woeginger’s hardness result [Woe97] for Vector Bin Packing can be easily modified to give APX-hardness for 2-dimensional Vector Bin Covering as well.

1.1 Our Results

We prove almost optimal hardness results for the three multidimensional problems discussed above.

1.1.1 Vector Bin Packing

For the Vector Bin Packing problem, we prove a $\Omega(\log d)$ asymptotic hardness of approximation when d is a large constant, matching the $\ln d + O(1)$ approximation algorithms [CK04, BCS09, BEK16], up to constants.

Theorem 1.1. *There exists an integer d_0 and a constant $c > 0$ such that for all constants $d \geq d_0$, d -dimensional Vector Bin Packing has no asymptotic $c \log d$ factor polynomial time approximation algorithm unless $\mathbf{P} = \mathbf{NP}$.*

We obtain our hardness result via a reduction from the set cover problem on certain structured instances. In the set cover problem, we are given a set system $\mathcal{S} \subseteq 2^V$ on a universe V , and the goal is to pick the minimum number of sets from \mathcal{S} whose union is V . Observe that Vector Bin Packing is a special case of the set cover problem with the vectors being the elements and every maximal set of vectors whose sum is at most 1 in every coordinate (known as “configurations”) being the sets. In fact, in the elegant Round & Approx framework [BCS09, BEK16], the Vector Bin Packing problem is viewed as a set cover instance, and the algorithms proceed by rounding the standard set cover LP. Towards proving the hardness of Vector Bin Packing, we ask the converse: *Which families of set cover instances can be cast as d -dimensional Vector Bin Packing?*

We formalize this question using the notion of *packing dimension* of a set system \mathcal{S} on a universe V : it is the smallest integer d such that there is an embedding $f : V \rightarrow [0, 1]^d$ such that a set $S \subseteq V$ is in \mathcal{S} if and only if

$$\left\| \sum_{v \in S} f(v) \right\|_{\infty} \leq 1$$

If a set system has packing dimension d , then the corresponding set cover problem can be embedded as a d -dimensional Vector Bin Packing instance. However, it is not clear if the hard instances of the set cover problem have a low packing dimension. Indeed the instances in the $(1 - \epsilon) \ln n$ set cover hardness [Fei98] have large packing dimension that grows with n , which we cannot afford as we are operating in the constant d regime. We get around this by starting our reduction from highly structured yet hard instances of set cover. In particular, we study simple bounded set systems which satisfy the following three properties:

1. The set system is *simple*⁴ i.e., every pair of sets intersect in at most one element.

⁴Simple set families are also known as linear set families.

2. The cardinality of each set is at most k , a fixed constant.
3. Each element in the family is present in at most $\Delta = k^{O(1)}$ sets.

Kumar, Arya, and Ramesh [KAR00] proved that simple set cover i.e., set cover with the restriction that every pair of sets intersect in at most one element, is hard to approximate within $\Omega(\log n)$. We observe that by modifying the parameters slightly in their proof, we can obtain the $\Omega(\log k)$ hardness of simple bounded set cover.

We prove that simple bounded set systems have packing dimension at most $k^{O(1)}$. Thus, the $\Omega(\log k)$ simple bounded set cover hardness translates to $\Omega(\log d)$ hardness of Vector Bin Packing when d is a constant. Note that the optimal value of the set cover instances can be made arbitrarily large in terms of k by starting with a Label Cover instance with an arbitrarily large number of edges. Thus, our Vector Bin Packing hardness holds for asymptotic approximation as well.

Our upper bound on the packing dimension is obtained in two steps: First, we write the given simple bounded set system as an intersection of $(k\Delta)^{O(1)}$ structured simple bounded set systems on the same universe, and then we give an embedding using $(k\Delta)^{O(1)}$ dimensions bounding the packing dimension of these structured simple bounded set systems. This idea of decomposition into structured instances is inspired from a work of Chandran, Francis, and Sivadasan [CFS08] where an upper bound on the Boxicity of a graph is obtained in terms of its maximum degree. We believe that the packing dimension of set systems is worth studying on its own, especially in light of its close connections to the notions of dimension of graphs such as Boxicity.

1.1.2 Vector Scheduling

For the Vector Scheduling problem, we obtain a $\Omega((\log d)^{1-\epsilon})$ hardness under $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$, almost matching the $O\left(\frac{\log d}{\log \log d}\right)$ algorithms [HS19, IKKP19].

Theorem 1.2. *For every constant $\epsilon > 0$, assuming $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$, d -dimensional Vector Scheduling has no polynomial time $\Omega((\log d)^{1-\epsilon})$ -factor approximation algorithm when d is part of the input.*

We obtain the hardness result via a reduction from the Monochromatic Clique problem. In the Monochromatic-Clique(k, B) problem, given a graph $G = (V, E)$ with $|V| = n$ and parameters $k(n)$ and $B(n)$, the goal is to distinguish between the case when G is k -colorable and the case when in any assignment of k -colors to vertices of G , there is a clique of size B all of whose vertices are assigned the same color. When $B = 2$, this is the standard graph coloring problem. Note that the problem gets easier as B increases. Indeed, when $B > \sqrt{n}$, we can solve the problem in polynomial time by computing the Lovász theta function of the complement graph. We are interested in proving the hardness of the problem for B as large a function of n as possible, for some k . For example, given a graph that is promised to be k colorable, can we prove the hardness of assigning k colors to the vertices of the graph in polynomial time where each color class has maximum clique at most $B = \log n$?

The Monochromatic Clique problem was defined formally by Im, Kell, Kulkarni, and Panigrahi [IKKP19] in the context of proving lower bounds for online Vector Scheduling. It was also used implicitly in the $\omega(1)$ NP-hardness of Vector Scheduling by Chekuri and Khanna [CK04]. They proved (implicitly) that Monochromatic Clique is NP-hard when B is an arbitrary constant using a reduction from $n^{1-\epsilon}$ hardness of graph coloring. We observe that the same reduction combined with better hardness of graph chromatic number [Kho01] proves the hardness of Monochromatic Clique when $B = (\log n)^\gamma$, for some constant $\gamma > 0$ under the assumption that $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$.

We then amplify this hardness to $B = (\log n)^C$, for every constant $C > 0$. Our main idea in this amplification procedure is the notion of a stronger form of Monochromatic Clique where given a graph and parameters k, B, C , the goal is to distinguish between the case that G is k colorable vs. in any k^C coloring of G , there is a monochromatic clique of size B . It turns out that the graph coloring hardness of Khot [Kho01] already proves the hardness of this stronger variant of Monochromatic clique when $B = (\log n)^\gamma$ for any constant C . We then use lexicographic product of graphs to amplify this result into the hardness of original Monochromatic Clique problem with $B = (\log n)^C$ for any constant C under the same assumption that $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$. This directly gives the required hardness of Vector Scheduling using the reduction in [CK04].

The Vector Scheduling problem is also closely related to the Balanced Hypergraph Coloring problem where the input is a hypergraph H and a parameter k , and the objective is to color the vertices of H using k colors to minimize the maximum number of times a color appears in an edge. We use this connection to improve upon the NP-hardness of the problem:

Theorem 1.3. *For every constant $C > 0$, d -dimensional Vector Scheduling is NP-hard to approximate within $\Omega((\log \log d)^C)$ when d is part of the input.*

Consider the case when each vector job is from $\{0, 1\}^d$. In this setting, we can view each coordinate as an edge in a hypergraph, and each vector corresponds to a vertex of the hypergraph. The goal is to find a m -coloring of vertices of the hypergraph i.e., an assignment of the vectors to m machines to minimize the maximum number of monochromatic vertices in an edge, which directly corresponds to the maximum load on a machine.

This problem of coloring a hypergraph to ensure that no color appears too many times in each edge is known as Balanced Hypergraph Coloring. Guruswami and Lee [GL18] obtained strong hardness results for this problem when k , the uniformity of the hypergraph is a constant, using the Label Cover Long Code framework combined with analytical techniques such as the invariance principle. However, when k is super constant, the invariance principle based methods give weak bounds as the soundness of the Label Cover has to be at least exponentially small in k . Recently, using combinatorial tools to analyze the gadgets instead of the standard analytical techniques, improvements have been obtained for various hypergraph coloring problems [Bha18, ABP19] in the super-constant inapproximability regime. We follow the same route and use combinatorial tools to analyze the gadgets in the Label Cover Long Code framework and obtain better hardness of Balanced Hypergraph Coloring in the regime of super-constant uniformity k . The key combinatorial lemma used in our analysis was proved recently by Austrin, Bhangale, and Potukuchi [ABP20] using a generalization of the Borsuk-Ulam theorem.

The NP-hardness of Vector Scheduling follows directly from the hardness of Balanced Hypergraph Coloring using the above-described reduction. This NP-hardness result uses near-linear size Label Cover hardness results [MR10, DS14]. By using the standard Label Cover hardness obtained by combining PCP Theorem and Parallel Repetition in the same reduction, we also prove an intermediate result bridging the above two hardness results for Vector Scheduling.

Theorem 1.4. *There exists a constant $\gamma > 0$ such that assuming $\text{NP} \not\subseteq \text{DTIME}\left(n^{O(\log \log n)}\right)$, d -dimensional Vector Scheduling is hard to approximate within $\Omega((\log d)^\gamma)$ when d is part of the input.*

1.1.3 Vector Bin Covering

For the Vector Bin Covering problem, we show $\Omega\left(\frac{\log d}{\log \log d}\right)$ hardness, almost matching the $O(\log d)$ factor algorithm [AAC⁺98].

Theorem 1.5. *d -dimensional Vector Bin Covering is NP-hard to approximate within $\Omega\left(\frac{\log d}{\log \log d}\right)$ factor when d is part of the input.*

Similar to Vector Scheduling, the hard instances for the Vector Bin Covering are when each vector is in $\{0, 1\}^d$. Using the same connection as before, we view this problem as a hypergraph coloring problem where each edge of the hypergraph corresponds to a coordinate, and each vertex corresponds to a vector. Assigning the vectors to the bins such that in each bin, the sum is at least 1 in every coordinate corresponds to assigning colors to vertices of the hypergraph such that in every edge, all the colors appear. Such a coloring of the hypergraph with k colors where all the k colors appear in every edge is known as a k -rainbow coloring of the hypergraph.

Strong hardness results are known for approximate rainbow coloring [GL18, ABP20, GS20] when k is a constant. While these results give decent bounds in the super constant regime, they proceed via the hardness of Label Cover whose soundness is an inverse polynomial function of k . Because of this, in the NP-hard instances, the number of edges is at least doubly exponential in k . Instead, by losing a factor of 2 in the hardness, we give a reduction to the approximate rainbow coloring problem from Label Cover with no gap i.e., just a “Label Coverized” 3-SAT instance. In these hard instances, the number of edges is single exponential in k , proving that it is NP-hard to distinguish the case that a hypergraph with m edges has a rainbow coloring with $\Omega\left(\frac{\log m}{\log \log m}\right)$ colors vs. it cannot be rainbow colored with 2 colors⁵. This hardness of approximate rainbow coloring gives the required Vector Bin Covering hardness immediately using the earlier mentioned analogy between rainbow coloring and Vector Bin Covering.

We summarize our results in Table 1.

1.2 Related Work

Online Algorithms. Multidimensional packing problems have been extensively studied in the online setting. For the d -dimensional Vector Bin Packing, the classical First-Fit algorithm [GGJY76] gives $O(d)$ competitive ratio, and Azar, Cohen, Kamara, and Shepherd [ACKS13] recently gave an almost matching $\Omega(d^{1-\epsilon})$ lower bound. For the d -dimensional Vector Scheduling, Im, Kell, Kulkarni and Panigrahi [IKKP19] gave a $O\left(\frac{\log d}{\log \log d}\right)$ competitive online algorithm and proved a matching lower bound. For the d -dimensional Vector Bin Covering problem, Alon et al. [AAC⁺98] gave a $2d$ competitive algorithm and proved a lower bound of $d + \frac{1}{2}$.

Geometric variants. There are various natural geometric variants of Vector Bin Packing that have been studied in the literature. A classical problem of this sort is the 2-dimensional Geometric Bin Packing, where the input is a set of rectangles that need to be packed into the minimum number of unit squares. After a long line of works, Bansal and Khan [BK14] gave a 1.405 factor asymptotic approximation algorithm for the problem. On the hardness front, Bansal and Sviridenko [BS04] showed that the problem does not admit an asymptotic PTAS, and this APX hardness result has been generalized to several related problems by Chlebík and Chlebíková [CC06]. We refer the reader to the excellent survey [CKPT17] regarding the geometric problems.

1.3 Organization

We first define the multidimensional problems and the Label Cover problem formally in Section 2. Next, we prove the hardness results for Vector Bin Packing, Vector Scheduling,

⁵Note that rainbow coloring with 2 colors is the same as proper 2 coloring (Property B) of the hypergraphs.

Problem	Subcase	Best Algorithm	Best Hardness
VBP	$d = 1$	PTAS [dVL81]	NP-Hard [GJ79]
	Fixed d	$\ln d + O(1)$ [BEK16]	$\Omega(\log d)$
	Arbitrary d	$1 + \epsilon d + O(\ln \frac{1}{\epsilon})$ [CK04]	$d^{1-\epsilon}$ [CK04, CKPT17]
VS	$d = 1$	EPTAS [Jan10]	No FPTAS [FKT89]
	Fixed d	PTAS [CK04]	No EPTAS [BOVdZ16]
	Arbitrary d	$O\left(\frac{\log d}{\log \log d}\right)$ [HS19, IKKP19]	$\Omega((\log d)^{1-\epsilon})$ $(\text{NP} \not\subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}}))$ $(\log d)^{\Omega(1)}$ $(\text{NP} \not\subseteq \text{DTIME}(n^{O(\log \log n)}))$ $(\log \log d)^{\omega(1)}$ (NP-Hardness)
VBC	$d = 1$	FPTAS [JS03]	NP-Hard [AJKL84]
	Arbitrary d	$O(\log d)$ [AAC+98]	$\Omega\left(\frac{\log d}{\log \log d}\right)$

Table 1: Approximation algorithms for the multidimensional packing problems. VBP, VS and VBC stand for Vector Bin Packing, Vector Scheduling and Vector Bin Covering respectively. All the results for VBP and the FPTAS for VBC are asymptotic. The results without citations are our new results.

and Vector Bin Covering in Section 3, Section 4 and Section 5 respectively. Finally, we conclude by mentioning a few open problems in Section 6.

2 Preliminaries

Notations. We use $[n]$ to denote the set $\{1, 2, \dots, n\}$. We use $\mathbf{1}^d$ to denote the d -dimensional vector $(1, 1, \dots, 1)$. For two d -dimensional real vectors \mathbf{a} and \mathbf{b} , we say that $\mathbf{a} \geq \mathbf{b}$ if $\mathbf{a}_i \geq \mathbf{b}_i$ for all $i \in [d]$. For a graph G , we let $\omega(G), \alpha(G), \chi(G)$ be the largest clique size, largest independent size, and the chromatic number of G respectively. A set system or set family \mathcal{S} on a universe V is a collection of subsets of V .

Problem Statements. We give formal definitions of the three problems that we study.

Definition 2.1. (*Vector Bin Packing*) In the Vector Bin Packing problem, the input is a set of n rational vectors $v_1, v_2, \dots, v_n \in [0, 1]^d$. The objective is to partition $[n]$ into minimum number of parts A_1, A_2, \dots, A_m such that

$$\left\| \sum_{j \in A_i} v_j \right\|_{\infty} \leq 1 \quad \forall i \in [m]$$

Definition 2.2. (*Vector Scheduling*) In the Vector Scheduling problem, the input is a set of n rational vector jobs $v_1, v_2, \dots, v_n \in [0, 1]^d$, and m identical machines. The objective is to assign the jobs to machines i.e. partition $[n]$ into m parts A_1, A_2, \dots, A_m to minimize the makespan which is defined as the maximum ℓ_{∞} load on a machine.

$$\max_{i \in [m]} \left\| \sum_{j \in A_i} v_j \right\|_{\infty}$$

Definition 2.3. (*Vector Bin Covering*) In the Vector Bin Covering problem, we are given n vectors $v_1, v_2, \dots, v_n \in [0, 1]^d$. The goal is to partition the input vectors into maximum number of parts A_1, A_2, \dots, A_m such that

$$\sum_{j \in A_i} v_j \geq \mathbf{1}^d \quad \forall i \in [m]$$

Asymptotic Approximation. For the Bin Packing problem, it is NP-Hard to identify if all the vectors can be packed into 2 bins or need 3 bins. This already proves that the problem is NP-hard to approximate within $\frac{3}{2}$ as per the usual notion of multiplicative approximation ratio. However, this is less interesting as there are much better *asymptotic* approximation algorithms for the problem which get $(1 + \epsilon)$ -factor approximation when the optimal value is large enough, for every positive constant $\epsilon > 0$.

Even for the Vector Bin Packing problem, the performance of an algorithm is typically measured in the asymptotic setting. We give the formal definition [CKPT17] of asymptotic approximation ratio of an algorithm \mathcal{A} for the Vector Bin Packing problem.

Definition 2.4. (*Asymptotic Approximation Ratio*) The asymptotic approximation ratio $\rho_{\mathcal{A}}^{\infty}$ of an algorithm \mathcal{A} for the Vector Bin Packing problem is

$$\rho_{\mathcal{A}}^{\infty} = \limsup_{n \rightarrow \infty} \rho_{\mathcal{A}}^n, \quad \rho_{\mathcal{A}}^n = \sup_{I \in \mathcal{I}} \left\{ \frac{\mathcal{A}(I)}{\text{OPT}(I)} : \text{OPT}(I) = n \right\}$$

where \mathcal{I} denotes the set of all possible Vector Bin Packing instances.

All the results mentioned in this paper regarding Vector Bin Packing are with respect to the asymptotic approximation ratio.

Label Cover. We define the Label Cover problem:

Definition 2.5. (*Label Cover*) In an instance of the Label Cover problem $G = (V = L \cup R, E, \Sigma_L, \Sigma_R, \Pi)$ with $|\Sigma_L| \geq |\Sigma_R|$, the input is a bipartite graph $L \cup R$ with constraints on every edge. The constraint on an edge e is a projection $\Pi_e : \Sigma_L \rightarrow \Sigma_R$. We say a labeling $\sigma : V \rightarrow \Sigma_L \cup \Sigma_R$ satisfies the constraint on the edge $e = (u, v)$ if $\Pi_e(\sigma(u)) = \sigma(v)$. The objective is to find a labeling $\sigma : V \rightarrow \Sigma_L \cup \Sigma_R$ that satisfies as many constraints as possible.

By a simple reduction from the 3-SAT problem, we can prove that Label Cover is NP-hard when Σ_L and Σ_R are constants (See e.g., Lemma 4.2 in [BG16]).

Theorem 2.6. Given a Label Cover instance when $\Sigma_L = \Sigma_R = [6]$, it is NP-hard to identify if it has a labeling that satisfies all the constraints.

The real use of Label Cover, however, lies in its strong hardness of approximation. PCP Theorem [ALM⁺98] combined with Raz's parallel repetition [Raz98] yields the following strong inapproximability of Label Cover problem.

Theorem 2.7. There exists an absolute constant $c > 1$ such that for every integer n and $\epsilon > 0$, there is a reduction from 3-SAT instance I over n variables to Label Cover instance $G = (V = L \cup R, E, \Sigma_L, \Sigma_R, \Pi)$ with $|V| \leq n^{O(\log(\frac{1}{\epsilon}))}$, $|\Sigma_L| \leq (\frac{1}{\epsilon})^c$ satisfying the following:

1. (*Completeness.*) If I is satisfiable, there exists a labeling to G that satisfies all the constraints.
2. (*Soundness.*) If I is not satisfiable, no labeling can satisfy an ϵ fraction of the constraints of G .

3. (Biregularity.) The graph $L \cup R, E$ is biregular with degrees on either side bounded by $\text{poly}\left(\frac{1}{\epsilon}\right)$.

Furthermore, the running time of the reduction is $\text{poly}\left(n, \frac{1}{\epsilon}\right)$.

Moshkovitz-Raz [MR10] proved the following hardness of near linear size Label Cover.

Theorem 2.8. *There exist absolute constants $c, c' > 1$ such that for every n and $\epsilon > 0$, there is a reduction from 3-SAT instance I over n variables to Label Cover instance $G = (V = L \cup R, E, \Sigma_L, \Sigma_R, \Pi)$ with $|V| \leq n^{1+o(1)} \left(\frac{1}{\epsilon}\right)^c$, $|\Sigma_L| \leq 2^{\left(\frac{1}{\epsilon}\right)^{c'}}$ satisfying the following:*

1. (Completeness.) *If I is satisfiable, there exists a labeling to G that satisfies all the constraints.*
2. (Soundness.) *If I is not satisfiable, no labeling can satisfy an ϵ fraction of the constraints of G .*
3. (Biregularity.) *The graph $L \cup R, E$ is biregular with degrees on either side $\text{poly}\left(\frac{1}{\epsilon}\right)$.*

Furthermore, when ϵ is a constant, the running time of the reduction is $\text{poly}(n)$.

3 Vector Bin Packing

In this section, we prove the hardness of approximation of Vector Bin Packing. First, we define the packing dimension of a set family and bound the packing dimension of simple set families. Next, we combine this upper bound with the hardness of set cover on simple bounded set systems to prove Theorem 1.1.

3.1 Packing Dimension

For a set family \mathcal{S} on a universe V , we define the packing dimension $\text{pdim}(\mathcal{S})$ below. For a function $f : V \rightarrow [0, 1]^K$ and a set $S \subseteq V$, we let $f(S)$ denote the vector $f(S) = \sum_{v \in S} f(v)$.

Definition 3.1. *For a set family \mathcal{S} on a universe V , the packing dimension $\text{pdim}(\mathcal{S})$ is defined as the smallest positive integer K such that there exists an embedding $f : V \rightarrow [0, 1]^K$ that satisfies the following property: For every set $S \subseteq V$, S is in the family \mathcal{S} if and only if*

$$\|f(S)\|_{\infty} \leq 1.$$

If no such embedding exists, we say that $\text{pdim}(\mathcal{S})$ is infinite.

For a set family \mathcal{S} to have finite packing dimension i.e. for an embedding $f : V \rightarrow [0, 1]^K$ realizing the above condition to exist requires two conditions:

1. The set family is *downward closed* i.e. for every $S \in \mathcal{S}$ and $T \subseteq S$, $T \in \mathcal{S}$ as well.
2. For every element $v \in V$, there is a set $S \in \mathcal{S}$ with $v \in S$. We call a set family \mathcal{S} on a universe V non-trivial if for every $v \in V$, there is a set $S \in \mathcal{S}$ with $v \in S$.

On the other hand, any set system that satisfies the above two conditions i.e. being downward closed and non-trivial has a finite packing dimension. Before proving this statement, we first prove the following simple but useful proposition.

Proposition 3.2. *For a pair of set families \mathcal{S}_1 and \mathcal{S}_2 defined on the same universe V such that $\text{pdim}(\mathcal{S}_1)$ and $\text{pdim}(\mathcal{S}_2)$ are finite,*

$$\text{pdim}(\mathcal{S}_1 \cap \mathcal{S}_2) \leq \text{pdim}(\mathcal{S}_1) + \text{pdim}(\mathcal{S}_2)$$

Proof. Let $K_1 = \text{pdim}(\mathcal{S}_1)$ and $K_2 = \text{pdim}(\mathcal{S}_2)$. Suppose that $f_1 : V \rightarrow [0, 1]^{K_1}$ be such that for every set $S \subseteq V$,

$$\|f_1(S)\|_\infty \leq 1$$

if and only if $S \in \mathcal{S}_1$. Similarly, let $f_2 : V \rightarrow [0, 1]^{K_2}$ be such that for every set $S \subseteq V$,

$$\|f_2(S)\|_\infty \leq 1$$

if and only if $S \in \mathcal{S}_2$. Consider the function $f : V \rightarrow [0, 1]^{K_1+K_2}$ defined as $f(v) = (f_1(v), f_2(v))$. Then, for every set $S \subseteq V$, $\|f(S)\|_\infty \leq 1$ if and only if $\|f_1(S)\|_\infty \leq 1$ and $\|f_2(S)\|_\infty \leq 1$. Thus, for every set $S \subseteq V$, $\|f(S)\|_\infty \leq 1$ if and only if $S \in \mathcal{S}_1$ and $S \in \mathcal{S}_2$, or equivalently, if $S \in \mathcal{S}_1 \cap \mathcal{S}_2$. Hence, the packing dimension of $\mathcal{S}_1 \cap \mathcal{S}_2$ is at most $K_1 + K_2$. \square

For a set $S \subseteq V$, let S^\uparrow be the family of sets $T \subseteq V$ such that $S \subseteq T$. Similarly, let S^\downarrow be the family of sets $T \subseteq V$ such that $T \subseteq S$. For a set system \mathcal{S} , we let \mathcal{S}^\uparrow (resp. \mathcal{S}^\downarrow) denote the union of S^\uparrow (resp. S^\downarrow) over all $S \in \mathcal{S}$.

Consider a set $S \subseteq V$ with $|S| > 1$. For the set family $2^V \setminus S^\uparrow$, we have the embedding $f : V \rightarrow [0, 1]$ defined as

$$f(v) = \begin{cases} \frac{1}{|S|} + \frac{1}{|S|^2}, & \text{if } v \in S \\ 0 & \text{otherwise.} \end{cases}$$

This shows that $\text{pdim}(2^V \setminus S^\uparrow) \leq 1$ for all $S \subseteq V$ with $|S| > 1$. Note that we have

$$\mathcal{S} = \bigcap_{S \notin \mathcal{S}} 2^V \setminus S^\uparrow$$

for every downward closed set system \mathcal{S} . Combined with Proposition 3.2, we obtain that for every non-trivial downward closed family \mathcal{S} on a universe V , $\text{pdim}(\mathcal{S}) \leq 2^{|V|}$.

We are interested in the classes of set families for which there is an efficient embedding with packing dimension being independent of $|V|$. In particular, the class of set families that we study are bounded set families where each set has cardinality at most k , and each element appears in at most Δ sets. We can show that such bounded set families that are downward closed and non-trivial have packing dimension at most $(k\Delta)^{O(\Delta)}$. Together with the $\Omega(\log k)$ hardness [Tre01] of k -set cover where each set has cardinality at most k , a fixed constant and each element appearing in $(\log k)^{O(1)}$ sets, this packing dimension bound gives the hardness of $(\log d)^{\Omega(1)}$ for the Vector Bin Packing problem when d is a large constant. Unfortunately, the exponential dependence on Δ is necessary for the packing dimension of bounded set systems, and thus, this approach does not yield the optimal $\Omega(\log d)$ hardness of Vector Bin Packing.

Instead of using arbitrary bounded set families, we bypass this barrier by using simple bounded set families. Recall that a set family is called simple if any two distinct sets in the family intersect in at most one element. It turns out that for simple bounded set families i.e. simple set families \mathcal{S} where each set has cardinality at most k , and each element appears in at most Δ sets, the packing dimension of \mathcal{S}^\downarrow can be upper bounded by $(k\Delta)^{O(1)}$. Together with the $\Omega(\log k)$ hardness of simple k -set cover (proved in Appendix A), we get the optimal $\Omega(\log d)$ hardness of Vector Bin Packing when d is a large constant. In the next subsection, we prove the packing dimension upper bound, and we use this upper bound to prove the hardness of Vector Bin Packing in Section 3.3.

3.2 Packing Dimension of Simple Bounded Set Families

The main embedding result that we prove is that the downward closure of simple set systems where each set has cardinality k and each element appears in at most Δ sets has packing dimension at most polynomial in k, Δ .

Theorem 3.3. *Suppose that \mathcal{S} is a simple non-trivial set system on a universe V where each set has cardinality at most $k \geq 2$ and each element appears in at most Δ sets. Then,*

$$\text{pdim}(\mathcal{S}^\downarrow) \leq (k\Delta)^{O(1)}$$

Furthermore, an embedding realizing the above can be found in time polynomial in $|V|$.

We prove the embedding result by writing the set family \mathcal{S}^\downarrow as an intersection of $(k\Delta)^{O(1)}$ structured set families each of which has packing dimension at most $(k\Delta)^{O(1)}$. We can then upper bound the packing dimension of \mathcal{S}^\downarrow using Proposition 3.2. The structured set systems we study are *sunflower-bouquets*, which are a disjoint union of sunflowers that have a single element as the kernel. The formal definition of the sunflower-bouquet set families is below. See Figure 1 for an illustration.

Definition 3.4. (*Sunflower-bouquets*) *A simple set system \mathcal{S} on a universe V is called a sunflower-bouquet with core $U \subseteq V, U \neq \emptyset$ if the following hold.*

1. *Every set $S \in \mathcal{S}$ satisfies $|S \cap U| = 1$. Furthermore, for every $u \in U$, there is a set $S \in \mathcal{S}$ with $u \in S$.*
2. *For any pair of sets $S_1, S_2 \in \mathcal{S}$ with $S_1 \cap S_2 \neq \emptyset$, we have $S_1 \cap U = S_2 \cap U = S_1 \cap S_2$.*

We now give an efficient embedding for a sunflower-bouquet \mathcal{S} on a universe V with core $U \subseteq V, U \neq \emptyset$. The main motivation behind this lemma is to upper bound the packing dimension of the set system $\mathcal{T}^\downarrow = \mathcal{S}^\downarrow \cup \{S \subseteq V \setminus U : |S| \leq k\}$.

Lemma 3.5. *Fix an integer $k \geq 2$. Let \mathcal{S} be a simple set family defined on a universe V that is a sunflower-bouquet with core U . Furthermore, each set in the family has cardinality at most k and each element appears in at most Δ sets. Then, there exists an embedding $f : V \rightarrow [0, 1]^K$ that satisfies*

(A) *For every set $S \in \mathcal{S}$,*

$$\|f(S)\|_\infty \leq 1.$$

(B) *For every set $S \notin \mathcal{S}^\downarrow$ with $S \cap U \neq \emptyset$,*

$$\|f(S)\|_\infty > 1.$$

(C) *For every set $S \subseteq V$ with $S \cap U = \emptyset$ and $|S| \leq k$,*

$$\|f(S)\|_\infty \leq 1.$$

(D) *For every set $S \subseteq V$ with $|S| > k$,*

$$\|f(S)\|_\infty > 1.$$

with $K = (k\Delta)^{O(1)}$. Furthermore, such an embedding can be found in time polynomial in $|V|$ given \mathcal{S} .

Proof. Let $U = \{u_1, u_2, \dots, u_m\}$. We can partition $V \setminus U$ into V_0, V_1, \dots, V_m with

$$V_i = \bigcup_{S \in \mathcal{S}: u_i \in S} S \setminus \{u_i\}$$

for all $i \in [m]$. As each set in \mathcal{S} has cardinality at most k and each element appears in at most Δ sets, we get that $|V_i| \leq k\Delta$ for all $i \in [m]$. For every $i \in [m]$, we order the elements of V_i as $\{v_{i,1}, v_{i,2}, \dots, v_{i,k\Delta}\}$ (with repetitions if needed).

We construct the embedding f in two steps.

Step-1. Eliminating the cross-sunflower sets. In the first step, our goal is to find an embedding $g : V \rightarrow [0, 1]^{K_1}$ with $K_1 = 2k\Delta + 2$ such that




Figure 1: An illustration of a sunflower-bouquet set family. Here, \mathcal{S} is the family of all the green colored sets. It is a sunflower-bouquet with core $U = \{u_1, u_2, u_3\}$. In the embedding, we ensure that the ℓ_∞ norm of the left red set is greater than 1 in the first step while the right side red set is handled in the second step.

1. g satisfies the conditions (A) and (C) i.e. for every set $S \in \mathcal{S}$, $\|g(S)\|_\infty \leq 1$, and for every set $S \subseteq V$ with $S \cap U = \phi$ and $|S| \leq k$, $\|g(S)\|_\infty \leq 1$.
2. For every ‘‘cross-sunflower’’ set $S \subseteq V$ with $u_i \in S$ for some $i \in [m]$, and $S \cap V_{i'} \neq \phi$ for $i' \in [m], i' \neq i$, we have $\|g(S)\|_\infty > 1$.

We achieve this by setting $g = (f_0, f_1, \dots, f_{k\Delta})$, where each $f_l : V \rightarrow [0, 1]^2, l \in \{0, 1, \dots, k\Delta\}$ satisfies the conditions (A) and (C), and overall, the embedding g satisfies the second condition above.

We define the embedding $f_0 : V \rightarrow [0, 1]^2$ as

$$f_0(v) = \begin{cases} (1, \frac{1}{k}), & \text{if } v \in U \\ (\frac{1}{k}, \frac{1}{k}) & \text{if } v \in V_0 \\ (0, \frac{1}{k}) & \text{otherwise.} \end{cases}$$

We can verify that f_0 satisfies the conditions (A) and (C).

We choose m distinct rational numbers $\alpha_1, \dots, \alpha_m$ with $1 - \frac{1}{k} < \alpha_i < 1$ for all $i \in [m]$. We define the embeddings $f_l : V \rightarrow [0, 1]^2, l \in [k\Delta]$ as follows. Consider an $l \in [k\Delta]$.

1. For $i \in [m]$, we set

$$f_l(u_i) = \left(\alpha_i, 2 - \frac{1}{k} - \alpha_i \right)$$

2. For $i \in [m]$ and $v_{i,j} \in V_i$, we set $f_l(v_{i,j}) = (0, 0)$ if $v_{i,j} \neq v_{i,l}$. We set

$$f_l(v_{i,l}) = \left(1 - \alpha_i, \alpha_i + \frac{1}{k} - 1 \right)$$

3. For $v \in V_0$, we set $f_l(v) = (0, 0)$.

We verify that these embeddings satisfy the conditions (A) and (C). Fix an $l \in [k\Delta]$.

(A) Consider a set $S \in \mathcal{S}$. Let $i \in [m]$ be such that $\{u_i\} = S \cap U$. We have

$$\begin{aligned} f_l(S) &= \sum_{v \in S} f_l(v) \\ &\leq \sum_{v \in \{u_i\} \cup V_i} f_l(v) \\ &= f_l(u_i) + f_l(v_{i,l}) \\ &= \left(\alpha_i, 2 - \frac{1}{k} - \alpha_i \right) + \left(1 - \alpha_i, \alpha_i + \frac{1}{k} - 1 \right) = (1, 1). \end{aligned}$$

(C) This follows directly from the fact that $\|f_l(v)\|_1 \leq \frac{1}{k}$ for all $l \in [k\Delta]$ and $v \in V \setminus U$.

Let $g : V \rightarrow [0, 1]^{2k\Delta+2}$ be defined as $g = (f_0, f_1, \dots, f_{k\Delta})$. As each of the individual embeddings satisfies (A) and (C), g also satisfies the conditions (A) and (C).

Let $S \subseteq V$ be such that

$$\|g(S)\|_\infty \leq 1.$$

i.e. $\|f_l(S)\|_\infty \leq 1$ for all $l \in \{0, 1, \dots, k\Delta\}$. We collect the following observations that will be used later when we analyze the final embedding.

1. As $f_0(v)_1 = 1$ for all $v \in U$, $|S \cap U| \leq 1$. As $f_0(v)_1 = \frac{1}{k}$ for all $v \in V_0$, if $S \cap U \neq \emptyset$, then $S \cap V_0 = \emptyset$.
2. As $f_0(v)_2 = \frac{1}{k}$ for all $v \in V$, $|S| \leq k$.
3. Suppose that $S \cap U = \{u_i\}$. Then, we claim that $S \subseteq \{u_i\} \cup V_i$. Suppose for contradiction that this is not the case, and there exists $v_{i',l} \in V_{i'}$ with $i' \neq i$, $i' \in [m]$ and $l \in [k\Delta]$ such that $v_{i',l} \in S$. We have

$$\begin{aligned} f_l(S) &= \sum_{v \in S} f_l(v) \\ &\geq f_l(u_i) + f_l(v_{i',l}) \\ &= \left(\alpha_i, 2 - \frac{1}{k} - \alpha_i \right) + \left(1 - \alpha_{i'}, \alpha_{i'} + \frac{1}{k} - 1 \right) \\ &= (1 + \alpha_i - \alpha_{i'}, 1 + \alpha_{i'} - \alpha_i) \end{aligned}$$

As $\alpha_i \neq \alpha_{i'}$, $\|f_l(S)\|_\infty > 1$, a contradiction.

Step 2. Pinning down the intra-sunflower sets. In the second step, our goal is to find an embedding $g' : V \rightarrow [0, 1]^{K_2}$ with $K_2 = (k\Delta)^2$ such that

1. g' satisfies the conditions (A) and (C).
2. For every $i \in [m]$ and “intra-sunflower” set $S \subseteq \{u_i\} \cup V_i$ such that $u_i \in S$ and $S \notin \mathcal{S}^\perp$, we have $\|g'(S)\|_\infty > 1$.

We achieve this by setting $g' = (g_1, g_2, \dots, g_{(k\Delta)^2})$ where each g_l , $l \in [(k\Delta)^2]$ satisfies the conditions (A) and (C), and the overall function g' satisfies the second condition above.

For every $i \in [m]$, we order all the pairs of distinct elements $x, y \in V_i$ as $\{V_{i,1}, V_{i,2}, \dots, V_{i,(k\Delta)^2}\}$ (with repetitions if needed). The upper bound on the number of such pairs is obtained using the fact that $|V_i| \leq k\Delta$ for all $i \in [m]$.

We define the embeddings $g_l : V \rightarrow [0, 1]$, $l \in [(k\Delta)^2]$ below. Fix an $l \in [(k\Delta)^2]$.

1. Consider an $i \in [m]$. We have two different cases:

- (a) If $V_{i,l} \cup \{u_i\} \in \mathcal{S}^\downarrow$, we set $g_l(u_i) = 0$ and $g_l(v) = 0$ for all $v \in V_i$.
- (b) If $V_{i,l} \cup \{u_i\} \notin \mathcal{S}^\downarrow$, we set $g_l(v) = \frac{1}{k}$ for all $v \in V_{i,l}$, and $g_l(v) = 0$ for all $v \in V_i \setminus V_{i,l}$. We set

$$g_l(u_i) = 1 - \frac{2}{k} + \frac{1}{k^2}$$

2. For all $v \in V_0$, we set $g_l(v) = 0$.

We now verify that these embeddings satisfy the conditions (A) and (C). Fix an integer $l \in [(k\Delta)^2]$.

- (A) Consider a set $S \in \mathcal{S}$. Let $\{u_i\} = S \cap U$. If $\{u_i\} \cup V_{i,l} \in \mathcal{S}^\downarrow$, $g_l(v) = 0$ for all $v \in S$, and thus we have $|g_l(S)| \leq 1$. Now suppose that $\{u_i\} \cup V_{i,l} \notin \mathcal{S}^\downarrow$. This implies that $V_{i,l}$ is not a subset of S . As $|V_{i,l}| = 2$, $|V_{i,l} \cap S| \leq 1$. We get

$$\begin{aligned} \sum_{v \in S} g_l(v) &= g_l(u_i) + \sum_{v \in S \cap V_i} g_l(v) \\ &= g_l(u_i) + \sum_{v \in S \cap V_{i,l}} g_l(v) \\ &\leq g_l(u_i) + \frac{1}{k} \\ &= 1 - \frac{2}{k} + \frac{1}{k^2} + \frac{1}{k} \leq 1 \end{aligned}$$

- (C) This follows from the fact that $g_l(v) \leq \frac{1}{k}$ for all $v \in V \setminus U$.

Final embedding. We define the final embedding $f : V \rightarrow [0, 1]^{2k\Delta+2+(k\Delta)^2}$ as $f = (g, g_1, g_2, \dots, g_{(k\Delta)^2})$. As each of these embeddings satisfies the conditions (A) and (C), the final embedding f also satisfies the conditions (A) and (C).

Suppose that $\|f(S)\|_\infty \leq 1$ for a set $S \subseteq V$. Then, $\|g(S)\|_\infty \leq 1$ and $\|g_l(S)\|_\infty \leq 1$ for all $l \in [(k\Delta)^2]$. Condition (D) follows easily as $\|g(S)\|_\infty \leq 1$ implies that $|S| \leq k$.

We now return to condition (B). Suppose that $S \subseteq V$ with $S \cap U \neq \emptyset$ satisfies $\|f(S)\|_\infty \leq 1$. Our goal is to show that $S \in \mathcal{S}^\downarrow$. We have already deduced from $\|g(S)\|_\infty \leq 1$ that $|S \cap U| = 1$ and $|S| \leq k$. Let $S \cap U = \{u_i\}$ and we also have $S \subseteq \{u_i\} \cup V_i$. Let $S = \{u_i, s_1, s_2, \dots, s_p\}$ where $s_j \in V_i$ for all $j \in [p]$. Note that for every $v \in V_i$, there is exactly one set $S(v) \in \mathcal{S}$ such that $v \in S(v)$ and this set $S(v)$ satisfies $u_i \in S(v)$. This follows from the definition of V_i and the fact that the set family \mathcal{S} is a sunflower-bouquet.

We now claim that $S(s_{j_1}) = S(s_{j_2})$ for all $j_1, j_2 \in [p]$. Suppose for contradiction that there exist $j_1, j_2 \in [p]$ with $S(s_{j_1}) \neq S(s_{j_2})$. This implies that $\{u_i, s_{j_1}, s_{j_2}\} \notin \mathcal{S}^\downarrow$ as otherwise, if there exists $T \in \mathcal{S}$ such that $\{u_i, s_{j_1}, s_{j_2}\} \subseteq T$, we have $S(s_{j_1}) = S(s_{j_2}) = T$. Let $l \in [(k\Delta)^2]$ be such that $V_{i,l} = \{s_{j_1}, s_{j_2}\}$. As $V_{i,l} \cup \{u_i\} \notin \mathcal{S}^\downarrow$, we have $g_l(v) = \frac{1}{k}$ for all $v \in V_{i,l}$ and

$$g_l(u_i) = 1 - \frac{2}{k} + \frac{1}{k^2}$$

Thus, we get that

$$\begin{aligned} \sum_{v \in S} g_l(v) &= g_l(u_i) + \sum_{v \in S \setminus \{u_i\}} g_l(v) \\ &= g_l(u_i) + \sum_{v \in V_{i,l}} g_l(v) \\ &= 1 - \frac{2}{k} + \frac{1}{k^2} + \frac{2}{k} = 1 + \frac{1}{k^2} \end{aligned}$$

contradicting the fact that $g_l(S) \leq 1$. This completes the proof that $S(s_{j_1}) = S(s_{j_2})$ for all $j_1, j_2 \in [p]$. Thus, there exists a set $S(s_1) \in \mathcal{S}$ such that $S \subseteq S(s_1)$, which implies that $S \in \mathcal{S}^\downarrow$. Thus, every set $S \subseteq V$ with $\|f(S)\|_\infty \leq 1$ and $S \cap U \neq \emptyset$ satisfies $S \in \mathcal{S}^\downarrow$, which proves the condition (B).

Note that our construction is explicit, and we have a polynomial time algorithm to output the required embedding. The dimension of the embedding is $2k\Delta + 2 + (k\Delta)^2$, which is at most $(k\Delta)^{O(1)}$. \square

As a corollary, we bound the packing dimension of the set family

$$\mathcal{T}^\downarrow = \mathcal{S}^\downarrow \cup \{S \subseteq V \setminus U : |S| \leq k\}.$$

Corollary 3.6. *Suppose that \mathcal{T} is a set family defined on a universe V with*

$$\mathcal{T} = \mathcal{S} \cup \{S \subseteq V \setminus U : |S| \leq k\}$$

where $\mathcal{S} \subseteq 2^V$ is a sunflower-bouquet with core U . Furthermore, each set in \mathcal{S} has cardinality at most $k \geq 2$ and each element appears in at most Δ sets in \mathcal{S} . Then,

$$\text{pdim}(\mathcal{T}^\downarrow) \leq (k\Delta)^{O(1)}$$

Furthermore, an embedding realizing this packing dimension can be found in time polynomial in $|V|$ given \mathcal{S} .

Proof. As \mathcal{S} is a sunflower-bouquet, from Lemma 3.5, there exists an embedding $f : V \rightarrow [0, 1]^K$ that satisfies the conditions (A), (B), (C) and (D) with $K = (k\Delta)^{O(1)}$. Conditions (A) and (C) together imply that

$$\|f(S)\|_\infty \leq 1$$

for all $S \in \mathcal{T}$. Note that

$$\mathcal{T}^\downarrow = \mathcal{S}^\downarrow \cup \{S \subseteq V \setminus U : |S| \leq k\}.$$

Suppose that $S \subseteq V$ is a subset of V with $S \notin \mathcal{T}^\downarrow$. If $S \cap U = \emptyset$, then $|S| > k$, which implies that $\|f(S)\|_\infty > 1$ using condition (D). If $S \cap U \neq \emptyset$, then $S \notin \mathcal{S}^\downarrow$ which implies that $\|f(S)\|_\infty > 1$ using condition (B). Thus, $\|f(S)\|_\infty \leq 1$ if and only if $S \in \mathcal{T}^\downarrow$. \square

We are now ready to prove our main embedding result i.e. Theorem 3.3.

Proof of Theorem 3.3. We define a graph $G = (V, E)$ as follows: two elements $u, v \in V$ are adjacent in G if there exist sets $S_1, S_2 \in \mathcal{S}$ (not necessarily distinct) such that $u \in S_1, v \in S_2, S_1 \cap S_2 \neq \emptyset$. As the cardinality of each set in \mathcal{S} is at most k and each element of V is present in at most Δ sets, the maximum degree of a vertex in G can be bounded above as

$$\Delta(G) \leq k(k-1)\Delta^2$$

Thus, the chromatic number of G is at most $L = \chi(G) \leq k(k-1)\Delta^2 + 1 \leq k^2\Delta^2$. Using the greedy coloring algorithm, we can partition V into L non-empty parts U_1, U_2, \dots, U_L such that each U_j is a independent set in G . For every $j \in [L]$, as U_j is an independent set in G , we have

1. For every set $S \in \mathcal{S}$, $|S \cap U_j| \leq 1$.
2. Any two sets $S_1, S_2 \in \mathcal{S}$ with $S_1 \cap U_j \neq \emptyset, S_2 \cap U_j \neq \emptyset$ and $S_1 \cap S_2 \neq \emptyset$ satisfy $S_1 \cap U_j = S_2 \cap U_j = S_1 \cap S_2$.

We now define the set families $\mathcal{S}_1, \mathcal{S}_2, \dots, \mathcal{S}_L$ as follows:

$$\mathcal{S}_j = \{S \in \mathcal{S} : S \cap U_j \neq \emptyset\} \cup \{S \subseteq V \setminus U_j : |S| \leq k\}$$

We claim that $\bigcap_{j \in [L]} \mathcal{S}_j^\downarrow = \mathcal{S}^\downarrow$. First, consider an arbitrary set $S \in \mathcal{S}^\downarrow$ and an integer $j \in [L]$. As $|S| \leq k$, irrespective of S intersects U_j or not, $S \in \mathcal{S}_j^\downarrow$. Thus, $\mathcal{S}^\downarrow \subseteq \mathcal{S}_j^\downarrow$ for all $j \in [L]$. Consider a non-empty set $S \notin \mathcal{S}^\downarrow$. As U_1, U_2, \dots, U_L is a partition of V , there exists a $j \in [L]$ such that $S \cap U_j \neq \emptyset$. As $S \notin \mathcal{S}^\downarrow$, $S \notin \mathcal{S}_j^\downarrow$. This implies that

$$\bigcap_{j \in [L]} \mathcal{S}_j^\downarrow = \mathcal{S}^\downarrow$$

Using Proposition 3.2, in order to bound the packing dimension of \mathcal{S}^\downarrow , it suffices to bound the packing dimension of \mathcal{S}_j^\downarrow , $j \in [L]$.

Fix an integer $j \in [L]$ and consider the set family \mathcal{S}_j^\downarrow . It is defined on the universe V and there exists a non-empty subset $U_j \subseteq V$ such that

$$\mathcal{S}_j = \mathcal{S}'_j \cup \{S \subseteq V \setminus U_j : |S| \leq k\}$$

with

$$\mathcal{S}'_j = \{S \in \mathcal{S} : S \cap U_j \neq \emptyset\}.$$

Here, \mathcal{S}'_j is a simple set system which satisfies the following properties:

1. Each set in \mathcal{S}'_j has cardinality at most $k \geq 2$ and each element appears in at most Δ sets in \mathcal{S}'_j .
2. Every set $S \in \mathcal{S}'_j$ satisfies $|S \cap U_j| = 1$. As \mathcal{S} is non-trivial, for every $u \in U_j$, there exists a set $S \in \mathcal{S}'_j$ with $u \in S$.
3. For every pair of sets $S_1, S_2 \in \mathcal{S}'_j$ with $S_1 \cap S_2 \neq \emptyset$, $S_1 \cap U_j = S_2 \cap U_j = S_1 \cap S_2$.

In other words, the set family \mathcal{S}'_j is a sunflower-bouquet with core U_j . Using Corollary 3.6, we get that $\text{pdim}(\mathcal{S}'_j) \leq (k\Delta)^{O(1)}$ for all $j \in [L]$, which completes the proof. \square

3.3 Hardness of Vector Bin Packing

We show that for large enough constant d , Vector Bin Packing is hard to approximate within $\Omega(\log d)$. Our hardness is obtained via the hardness of set cover on simple bounded instances.

In the set cover problem, the input is a set family \mathcal{S} on a universe V with $|V| = n$. The objective is to pick the minimum number of sets $\{S_1, S_2, \dots, S_m\} \subseteq \mathcal{S}$ from the family such that their union is equal to V . The greedy algorithm where we repeatedly pick the set that covers the maximum number of new elements achieves a $\ln n$ approximation factor. Fiege [Fei98] proved a matching hardness of $(1 - \epsilon)(\ln n)$. On set systems where each pair of sets intersect in at most one element i.e. simple instances, $\Omega(\log n)$ hardness of set cover is proved by Kumar, Arya, and Ramesh [KAR00]. We observe that by changing the parameters slightly, their reduction also implies the same hardness on instances where the maximum set size is bounded:

Theorem 3.7. *(Set Cover on simple bounded instances) There exists an integer B_0 such that for every constant $B \geq B_0$, the Set Cover problem on simple set systems in which each set has cardinality at most B is NP-hard to approximate within $\Omega(\log B)$. Furthermore, in the hard instances, each element occurs in at most $O(B)$ sets.*

The details of the parameter modification appear in Appendix A.

We combine this set cover hardness with the bound on the packing dimension of simple set systems to prove the hardness of Vector Bin Packing.

Proof of Theorem 1.1. We prove the result by giving an approximation preserving reduction from the NP-hard problem of set cover on simple bounded set systems. Let \mathcal{S} be the set system from Theorem 3.7 defined on a universe V . Note that each set in the family has cardinality at most $k = B$ and each element in the universe appears in at most $\Delta = O(B)$ sets. We now output a set \mathcal{V} of $|V|^d$ vectors in $[0, 1]^d$ such that

1. (Completeness.) If there is a set cover of size m in \mathcal{S} , there is a packing of \mathcal{V} using m bins.
2. (Soundness.) If there is no set cover of size m' in \mathcal{S} , there is no packing of \mathcal{V} using m' bins.

We use Theorem 3.3 to compute an embedding $f : V \rightarrow [0, 1]^d$ in polynomial time such that

$$\|f(S)\|_\infty \leq 1$$

if and only if $S \in \mathcal{S}^\downarrow$, with $d = (k\Delta)^{O(1)} = B^{O(1)}$. Our output Vector Bin Packing instance is the set of vectors $f(v), v \in V$.

$$\mathcal{V} = \{f(v) : v \in V\}$$

Completeness. Suppose that there exist sets $S_1, S_2, \dots, S_m \in \mathcal{S}$ whose union is V . Then, we use m bins with the vectors $\{f(v_j) : j \in S_i\}$ in the i th bin. A vector might appear in multiple bins, but we can arbitrarily pick one bin for each vector while still maintaining the property that in each bin, the ℓ_∞ norm of the sum of the vectors is at most 1.

Soundness. Suppose that the minimum set cover in \mathcal{S} has cardinality at least $m' + 1$. Then, we claim that the set of vectors \mathcal{V} needs $m' + 1$ bins to be packed. Suppose for contradiction that there is a vector packing with m' bins. In other words, there exists a partition of V into $B_1, B_2, \dots, B_{m'}$ such that $\|f(B_i)\|_\infty \leq 1$ for all $i \in [m']$. As $\|f(B_i)\|_\infty \leq 1$, $B_i \in \mathcal{S}^\downarrow$ for all $i \in [m']$. That is, for every $i \in [m']$, there exists a set $S_i \in \mathcal{S}$ such that $B_i \subseteq S_i$. This implies that $\{S_1, S_2, \dots, S_{m'}\}$ is a set cover of V , a contradiction.

As the original bounded simple set cover problem is hard to approximate within $\Omega(\log B) = \Omega(\log d)$, the resulting Vector Bin Packing is hard to approximate within $\Omega(\log d)$. Furthermore, in the hard instances, the optimal value i.e. the minimum number of bins needed to pack the vectors can be made arbitrarily large, and thus, the hardness applies to the asymptotic approximation ratio. \square

4 Vector Scheduling

4.1 Monochromatic Clique

In the Monochromatic Clique problem, given a graph $G = ([n], E)$ and a parameter $k(n)$, the objective is to assign k colors to the vertices of G so as to minimize the largest monochromatic clique. More formally, we study the following decision version of the problem.

Definition 4.1. (*Monochromatic-Clique(k, B)*) In the Monochromatic-Clique(k, B) problem, given a graph $G = (V, E)$ with $|V| = n$ and parameters $k(n), B(n)$, the goal is to distinguish between the following:

1. (YES case) The chromatic number of G is at most k .
2. (NO case) In any assignment of k colors to the vertices of G , there is a clique of size B , all of whose vertices are assigned the same color.

It generalizes the standard k -Coloring problem, which corresponds to the case when $B = 2$. Note that the problem gets easier as B increases. Indeed, when $B > \sqrt{n}$, we can solve the problem in polynomial time using the canonical SDP relaxation. We present this algorithm and an almost matching integrality gap in Appendix B.

On the hardness front, we now prove that Monochromatic-Clique(k, B) is hard when $B = (\log n)^C$, for any constant C . We achieve this in two steps: First, we observe that the existing chromatic number hardness results already imply the hardness of monochromatic clique when $B = (\log n)^\gamma$ for some constant $\gamma > 0$. Next, we amplify this hardness by using lexicographic graph product.

4.1.1 Basic Hardness

We start with a couple of basic Ramsey theoretic lemmas from [CK04].

Lemma 4.2. For a graph $G = (V, E)$ with $|V| = n$, if $\omega(G) \leq B$, then $\alpha(G) \geq n^{\frac{1}{B}}$.

Lemma 4.3. For a graph $G = (V, E)$ with $|V| = n$, if $\omega(G) \leq B$, then $\chi(G) \leq O(n^{1-\frac{1}{B}} \log n)$.

We can use the above lemmas to prove that if the chromatic number of a graph is large enough, then in any assignment of k colors to the vertices of the graph, there is a large monochromatic clique.

Lemma 4.4. For every constant $\epsilon > 0$, if a graph $G = (V, E)$ with $|V| = n$ satisfies $\chi(G) \geq k \frac{n}{2^{(\log n)^\alpha}}$ for some integer k and $0 < \alpha < 1$, then in any assignment of k colors to V , there is a monochromatic clique of size $B = \Omega((\log n)^{1-\alpha-\epsilon})$.

Proof. Suppose for contradiction that there is an assignment of k colors to V without a monochromatic clique of size B . Using Lemma 4.3, the subgraphs corresponding to each of the k color classes has chromatic number at most

$$O(n^{1-\frac{1}{B}} \log n) = \frac{n}{2^{\Omega((\log n)^{\alpha+\epsilon})}} \log n < \frac{n}{2^{(\log n)^\alpha}}$$

colors. Thus, the whole graph has chromatic number at most $k \frac{n}{2^{(\log n)^\alpha}}$ colors, a contradiction. \square

Khot [Kho01] proved that assuming $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$, the chromatic number of graphs is hard to approximate within a factor of $\frac{n}{2^{(\log n)^{1-\gamma}}}$ for an absolute constant $\gamma > 0$. More formally, he proved the following:

Theorem 4.5. ([Kho01]) There exists a constant $\gamma > 0$, a function $k = k(n)$, and a randomized reduction that takes as input a 3-SAT instance I on n variables and outputs a graph $G = (V, E)$ with $|V| = N = 2^{\log n^{O(1)}}$ such that

1. (Completeness) If I is satisfiable, $\chi(G) \leq k$.
2. (Soundness) If I is not satisfiable, with probability at least $\frac{1}{2}$, $\chi(G) > k \frac{N}{2^{(\log N)^{1-\gamma}}}$.

Futhermore, the reduction runs in time $\text{poly}(N) = 2^{(\log n)^{O(1)}}$.

We observe that Khot's chromatic number hardness immediately gives $(\log n)^{\Omega(1)}$ hardness of Monochromatic Clique.

Lemma 4.6. *There exists a constant $\gamma > 0$, a function $k = k(n)$ such that the following holds. Assuming $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$, given a graph $G = ([n], E)$, there is no $n^{(\log n)^{O(1)}}$ time algorithm for $\text{Monochromatic-Clique}(k, B)$ when $B = \Omega((\log n)^\gamma)$.*

Proof. Using Khot's reduction, we get that there exists an absolute constant $\gamma > 0$ such that assuming $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$, given a graph $G = ([n], E)$ and a parameter $k(n)$, there is no $n^{(\log n)^{O(1)}}$ time algorithm to distinguish between the following:

1. (Completeness) $\chi(G) \leq k$.
2. (Soundness) $\chi(G) > k \frac{n}{2^{(\log n)^{1-\gamma}}}$.

Using Lemma 4.4, the Soundness condition implies that in any assignment of k colors to G , there is a monochromatic clique of size $\Omega((\log n)^{\gamma-\epsilon})$, for any constant $\epsilon > 0$. Thus, given a graph G and a parameter k , assuming $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$, there is no $n^{(\log n)^{O(1)}}$ time algorithm to distinguish between the following:

1. (Completeness) $\chi(G) \leq k$.
2. (Soundness) In any assignment of k colors to the vertices of G , there is a monochromatic clique of size $\Omega((\log n)^{\gamma'})$.

for any constant $\gamma' < \gamma$. □

4.1.2 Amplification using Lexicographic Product

We cannot directly amplify the hardness of the Monochromatic-Clique problem by taking graph products as we cannot preserve the chromatic number and also amplify the largest clique in an assignment of k colors at the same time. We get around this issue by defining a harder variant of Monochromatic Clique called Strong Monochromatic Clique and then amplifying it.

Definition 4.7. (*Strong Monochromatic-Clique(k, B, C)*) *In the Strong Monochromatic-Clique(k, B, C), given a graph G and parameters $k(n), B(n), C$, the goal is to distinguish between the following two cases:*

1. (YES case) *The chromatic number of G is at most k .*
2. (NO case) *In any assignment of k^C colors to the vertices of G , there is a monochromatic clique of size B .*

We now observe that the chromatic number hardness of Khot [Kho01] implies the same hardness as Lemma 4.6 for Strong Monochromatic Clique as well.

Lemma 4.8. *There exists a constant $\gamma > 0$ and a function $k = k(n)$ such that for every constant $C \geq 1$, the following holds. Assuming $\text{NP} \not\subseteq \text{ZPTIME}\left(n^{(\log n)^{O(1)}}\right)$, there is no $n^{(\log n)^{O(1)}}$ time algorithm for $\text{Strong Monochromatic-Clique}(k, B, C)$ when $B = \Omega((\log n)^\gamma)$.*

Proof. Note that the function k in Theorem 4.5 satisfies $k = o\left(2^{(\log N)^{1-\gamma}}\right)$. Thus, we can replace the soundness condition in Theorem 4.5 with $\chi(G) \geq k^C \frac{N}{2^{C(\log N)^{1-\gamma}}}$. Using Lemma 4.4, this implies that in any assignment of k^C colors to the vertices of G , there is a monochromatic clique of size $\Omega((\log N)^{\gamma-\epsilon})$, where $\epsilon > 0$ is an absolute constant. The hardness of Strong Monochromatic Clique then follows along the same lines as Lemma 4.6. \square

We amplify the hardness of **Strong Monochromatic-Clique**(k, B, C) to **Monochromatic-Clique**(k^C, B^C) using the lexicographic product of graphs. First, we define lexicographic product and prove some properties of it.

Definition 4.9. (*Lexicographic product of graphs*) Given two graphs G and H , the *Lexicographic graph product* $G \cdot H$ has vertex set $V(G) \times V(H)$, and two vertices $(u_1, v_1), (u_2, v_2)$ are adjacent if either $(u_1, u_2) \in E(G)$ or $u_1 = u_2$ and $(v_1, v_2) \in E(H)$.

The lexicographic product can be visualized as replacing each vertex of G with a copy of H and forming complete bipartite graphs between copies of vertices adjacent in G . For ease of notation, we let $G^2 = G \cdot G$. More generally, for an integer n that is a power of 2, we define G^n as taking the above lexicographic product of G with itself recursively $\log n$ times.

Lemma 4.10. *Let $n \geq 2$ be a power of 2. If $\chi(G) \leq k$, then $\chi(G^n) \leq k^n$.*

Proof. We prove that $\chi(G^2) \leq k^2$, and the statement follows by induction on n . If $f : G \rightarrow [k]$ is a proper k -coloring of G , then the coloring $f'(u, v) = (f(u), f(v))$ is a proper k^2 -coloring of $G \times G$. \square

Lemma 4.11. *Let $n \geq 2$ be a power of 2. Suppose that in any assignment of k colors to the vertices of G , there is a monochromatic clique of size B . Then, in any assignment of k colors to the vertices of G^n , there is a monochromatic clique of size B^n .*

Proof. We prove the statement for $n = 2$ and the lemma follows by induction on n . Let $f : V(G^2) \rightarrow [k]$ be a given assignment. For a vertex $v \in G$, consider the assignment $g_v : V(G) \rightarrow [k]$ defined as $g_v(u) = f(v, u)$. As every assignment of k colors to the vertices of G has a monochromatic clique of size B , there is a color $\alpha(v) \in [k]$ and a clique $S(v) \subseteq V(G)$ with $|S(v)| \geq B$ such that $g_v(u) = \alpha(v)$ for all $u \in S(v)$, or in other words, $f(v, u) = \alpha(v)$ for all $u \in S(v)$. Note that such a set $S(v)$ and $\alpha(v)$ exist for $v \in V(G)$. The function $\alpha : V(G) \rightarrow [k]$ can also be visualized as an assignment of k colors to the vertices of G , and thus there is a monochromatic clique T of size at least B with respect to this assignment. The set

$$\{S(v) : v \in T\}$$

is a monochromatic clique of size B^2 with respect to f in G . \square

By using the lexicographic product, we can get a polynomial time reduction from Strong Monochromatic Clique to Monochromatic Clique.

Lemma 4.12. *For every constant $C \geq 1$ that is a power of 2, there exists a polynomial time reduction from **Strong Monochromatic-Clique**(k, B, C) to **Monochromatic-Clique**(k^C, B^C).*

Proof. Given a graph G as an instance of **Strong Monochromatic-Clique**(k, B, C), we compute the graph $G' = G^C$. We claim that solving **Monochromatic-Clique**(k^C, B^C) on G' solves the original Strong Monochromatic Clique problem.

1. (Completeness.) Suppose that $\chi(G) \leq k$. Then, by Lemma 4.10, $\chi(G') \leq k^C$.

2. (Soundness.) Suppose that in any assignment of k^C colors to the vertices of G , there is a monochromatic clique of size B . Then, by Lemma 4.11, in any assignment of k^C colors to the vertices of G' , there is a monochromatic clique of size B^C . \square

Putting everything together, we obtain the following hardness of Monochromatic Clique.

Theorem 4.13. *For every constant $C > 0$, there exists a function $k = k(n)$ such that the following holds. Assuming $NP \not\subseteq ZPTIME\left(n^{(\log n)^{O(1)}}\right)$, there is no $n^{(\log n)^{O(1)}}$ time algorithm for Monochromatic-Clique(k, B) when $B = \Omega((\log n)^C)$.*

Proof. The proof follows directly by combining Lemma 4.8 and Lemma 4.12. \square

4.2 From Monochromatic Clique to Vector Scheduling

We now prove Theorem 1.2 using the above hardness of Monochromatic Clique.

Proof of Theorem 1.2. The reduction from Monochromatic-Clique(k, B) to Vector Scheduling is (implicitly) proved in [CK04]. We present it here for the sake of completeness. Given a graph $G = (V = [n], E)$, parameters k and B , we order all the B -sized cliques of G as T_1, T_2, \dots, T_d with $d \leq n^B$. We define a set of n vectors v_1, v_2, \dots, v_n of dimension d with

$$(v_i)_j = \begin{cases} 1 & \text{if } i \in T_j \\ 0 & \text{otherwise.} \end{cases}$$

The instance of the Vector Scheduling has these n vectors as the input and the number of machines is equal to k .

We analyze the reduction.

1. (Completeness.) Suppose that there exists a proper k -coloring of G , $c : V \rightarrow [k]$. We assign the vector v_i to the machine $c(i)$. For every $j \in [d]$, all the B vectors that have 1 in the j th dimension are assigned to distinct machines. Thus, the makespan of the scheduling is at most 1.
2. (Soundness.) Suppose that in any assignment of k colors to the vertices of G , there is a monochromatic clique of size B . In this case, the makespan of the scheduling is at least B .

We set $B = (\log n)^C$ for a large constant C to be set later. We choose k from Theorem 4.13 such that assuming $NP \not\subseteq ZPTIME\left(n^{(\log n)^{O(1)}}\right)$, there is no $n^{(\log n)^{O(1)}}$ time algorithm for Monochromatic-Clique(k, B). By the above reduction, we can conclude that there is no polynomial time algorithm that approximates the resulting Vector Scheduling instances within a factor of $B = (\log n)^C$. As $d \leq n^B$, we get that $\log d \leq (\log n)^{C+1}$, and $B \geq (\log d)^{1 - \frac{1}{C+1}}$. Setting $C = \frac{1}{\epsilon} - 1$, we get that d -dimensional Vector Scheduling has no polynomial time $\Omega((\log d)^{1-\epsilon})$ approximation algorithm assuming $NP \not\subseteq ZPTIME\left(n^{(\log n)^{O(1)}}\right)$, for every constant $\epsilon > 0$. \square

Remark 4.14. In [IKKP19], Im, Kell, Kulkarni, and Panigrahi also study the ℓ_r -norm minimization of Vector Scheduling where the objective is to minimize

$$\max_{k \in [d]} \left(\sum_{i=1}^m (L_i^k)^r \right)^{\frac{1}{r}}$$

where L_i^k denotes the load on the machine i on the k th dimension. They gave an algorithm with an approximation ratio $O\left(\left(\frac{\log d}{\log \log d}\right)^{1-\frac{1}{r}}\right)$. Our reduction from Monochromatic Clique gives almost optimal hardness for this variant as well: we get the hardness of $\Omega\left((\log d)^{1-\frac{1}{r}-\epsilon}\right)$ assuming $NP \not\subseteq ZPTIME\left(n^{(\log n)^{O(1)}}\right)$, for every constant $\epsilon > 0$.

4.3 Hardness of Vector Scheduling via Balanced Hypergraph Coloring

Observe that the resulting Vector Scheduling instances in the above reduction satisfy a stronger property: the vectors are from $\{0, 1\}^d$. In the setting where the vectors are from $\{0, 1\}^d$, the Vector Scheduling problem is closely related to the Balanced Hypergraph Coloring problem. In this problem, given a hypergraph H and an integer k , the objective is to assign k colors to the vertices of H minimizing the maximum number of monochromatic vertices in an edge. More formally, we study the following decision version of the problem.

Definition 4.15. (*Balanced Hypergraph Coloring.*) *In the Balanced Hypergraph Coloring problem, given a s -uniform hypergraph H and parameters k and $c < s$, the objective is to distinguish between the following:*

1. *There is an assignment of k colors to the vertices of H such that in every edge, each color appears at most c times.*
2. *The hypergraph H has no proper coloring with k colors i.e., in any assignment of k colors to the vertices of H , there is an edge all of whose s vertices are assigned the same color.*

We give a simple reduction from Balanced Hypergraph Coloring to Vector Scheduling.

Lemma 4.16. *Given a s -uniform hypergraph $H = (V' = [n'], E')$ and parameters k, c , there is a polynomial time reduction that outputs a Vector Scheduling instance I over n' vectors $v_1, v_2, \dots, v_{n'} \in \{0, 1\}^d$ on m' machines with $m' = k, d = |E'|$ such that*

1. *(Completeness.) If there is an assignment of k colors to the vertices of H such that each color appears at most c times in every edge, then there is a scheduling of I with makespan at most c .*
2. *(Soundness.) If H has no proper coloring with k colors, then in any scheduling of I , the makespan is at least s .*

Proof. Let $d = |E'|$. Order the edges of the hypergraph H as e_1, e_2, \dots, e_d . We define the set of vectors $v_1, v_2, \dots, v_{n'} \in \{0, 1\}^d$ as follows:

$$(v_i)_j = \begin{cases} 1 & \text{if } i \in e_j \\ 0 & \text{otherwise.} \end{cases}$$

We set the number of machines m' to be equal to the number of colors k . There is a natural correspondence between the assignment of k -colors to the vertices of H $f : V' \rightarrow [k]$, and the scheduling where we assign the vector v_i to the machine $f(i)$. We now analyze our reduction.

1. (Completeness.) If there exists an assignment of k colors $f : V' \rightarrow [k]$ where each color appears at most c times in each edge, we assign the vector $v_i, i \in [n']$ to the machine $f(i)$. In any dimension $j \in [d]$, at most c vectors v_i with $(v_i)_j = 1$ are scheduled on any machine. Thus, in any machine, the total load in each dimension is at most c .

2. (Soundness.) If there exists a vector scheduling $f : [n] \rightarrow [m']$ with makespan strictly smaller than s , assign the color $f(i)$ to the i th vertex of the hypergraph. In any edge of the hypergraph, each color appears fewer than s times as the makespan is smaller than s . Thus, $f : V' \rightarrow [k]$ is a proper k -coloring of the hypergraph H . \square

We prove the hardness results for Vector Scheduling, namely Theorem 1.3 and Theorem 1.4 by combining this reduction with the hardness of Balanced Hypergraph Coloring. Note that the dimension of the resulting instances in the above reduction is equal to m , the number of edges in the hypergraph H , and the ratio of the makespans in the completeness and soundness is equal to $\frac{s}{c}$. Thus, our goal is to prove the hardness of the Balanced Hypergraph Coloring problem where $\frac{s}{c}$ is as large as possible, as a function of m , the number of edges in the underlying hypergraph.

Towards this, we first give a reduction from the Label Cover problem to the Balanced Hypergraph Coloring problem.

Lemma 4.17. *Fix an odd prime number $k \geq 3$ and let $\epsilon = \frac{1}{k^8}$. Given a Label Cover instance $G = (V = L \cup R, E, \Sigma_L, \Sigma_R, \Pi)$, there is a polynomial time reduction that outputs a k^2 uniform hypergraph $H = (V', E')$ with $|V'| \leq |L|k^{|\Sigma_L|}$ such that*

1. (Completeness) *If G is satisfiable, there is an assignment of k colors to the vertices of H such that in every edge, each color occurs at most $2k$ times.*
2. (Soundness) *If no labeling to G can satisfy an ϵ fraction of the constraints, then H has no proper k -coloring, that is, in any assignment of k colors to the vertices of H , there is an edge all of whose vertices are assigned the same color.*

Furthermore, $|E'|$ is at most $|R|\Delta^k k^{|\Sigma_L|k^2}$ where Δ is the maximum degree of a vertex $v \in R$.

We defer the proof of Lemma 4.17 to Section 4.4.

Using Lemma 4.17, we can prove the hardness of Balanced Hypergraph Coloring via Label Cover hardness results. We obtain two different hardness results for the Balanced Hypergraph Coloring problem, one under $\text{NP} \not\subseteq \text{DTIME}(n^{O(\log \log n)})$ and another NP-hardness result, by using two different hardness results for the Label Cover problem. These two hardness results prove Theorem 1.4 and Theorem 1.3 respectively, using Lemma 4.16.

First, using the standard Label Cover hardness obtained using PCP Theorem [ALM⁺98] combined with Raz's Parallel Repetition theorem [Raz98], we get the following hardness of Balanced Hypergraph Coloring.

Theorem 4.18. *Assuming $\text{NP} \not\subseteq \text{DTIME}(n^{O(\log \log n)})$, there is no polynomial time algorithm for the following problem. Given a k^2 -uniform hypergraph $H = (V', E')$ with $m = |E'|$ and $k = (\log m)^{\Omega(1)}$, distinguish between the following:*

1. *There is an assignment of k colors to the vertices of H such that in any edge of the hypergraph, each color appears at most $2k$ times.*
2. *The hypergraph H has no proper k coloring.*

Proof. By setting $\epsilon = \frac{1}{k^8}$ in Theorem 2.7, we have a reduction from the 3-SAT problem on n variables to the Label Cover problem $G = (V = L \cup R, E, \Sigma_L, \Sigma_R, \Pi)$ with soundness ϵ and $|V| \leq n^{O(\log k)}$, $|\Sigma_L| \leq k^{O(1)}$ and $\Delta \leq k^{O(1)}$. Using Lemma 4.17, we can reduce this Label Cover instance to a Balanced Hypergraph Coloring instance $H = (V', E')$ with $|V'| \leq n^{O(\log k)} 2^{k^{O(1)}}$ and $|E'| \leq n^{O(\log k)} 2^{k^{O(1)}}$. We set $k = (\log n)^{\Omega(1)}$ such that $|V'| = n^{O(\log \log n)}$ and $|E'| = n^{O(\log \log n)}$ to obtain the required hardness of Balanced Hypergraph Coloring. \square

The proof of Theorem 1.4 follows immediately from Theorem 4.18 and Lemma 4.16.

Next, using the hardness of near linear sized Label Cover due to Moshkovitz and Raz [MR10], we obtain the following NP-hardness of Balanced Hypergraph Coloring.

Theorem 4.19. *For any constant $C \geq 1$, given a k^2 uniform hypergraph $H = (V', E')$ with $m = |E'|$ and $k = (\log \log m)^C$, it is NP-hard to distinguish between the following:*

1. *There is an assignment of k colors to the vertices of H such that in any edge of the hypergraph, each color appears at most $2k$ times.*
2. *The hypergraph H has no proper k coloring.*

Proof. By setting $\epsilon = \frac{1}{k^8}$ in Theorem 2.8, we can reduce a 3-SAT instance over n variables to a Label Cover instance $G = (V = L \cup R, E, \Sigma_L, \Sigma_R, \Pi)$ with soundness ϵ and $|V| \leq n^{1+o(1)}k^{O(1)}$, $|\Sigma_L| \leq 2^{k^{O(1)}}$, $\Delta = k^{O(1)}$. By using Lemma 4.17, we can reduce the Label Cover instance to a Balanced Hypergraph Coloring instance $H = (V', E')$ with $|V'| \leq n^{1+o(1)}2^{2k^{O(1)}}$ and $|E'|$ at most $n^{1+o(1)}2^{2k^{O(1)}}$. We set $k = (\log \log n)^{\Omega(1)}$ to obtain $|V'| = O(n^2)$, $|E'| = O(n^2)$.

Dinur and Steurer [DS14] gave an improvement to [MR10]—in the new Label Cover hardness, the alphabet size $|\Sigma_L|$ can be taken to be $2^{\left(\frac{1}{\epsilon}\right)^\gamma}$ for every constant $\gamma > 0$. Using this improved Label Cover hardness, we can set $k = (\log \log n)^C$ for any constant $C \geq 1$ in the hardness of Balanced Hypergraph Coloring. \square

The proof of Theorem 1.3 follows immediately from Theorem 4.19 and Lemma 4.16.

Finally, we remark that if the structured graph version of the Projection Games Conjecture [Mos15] holds, Lemma 4.17 and Lemma 4.16 together prove that d -dimensional Vector Scheduling is NP-hard to approximate within a factor of $(\log d)^{\Omega(1)}$.

4.4 Proof of Lemma 4.17

We follow the standard Label Cover-Long Code framework—see e.g., [ABP20].

Reduction. For ease of notation, let $n = |\Sigma_L|$. For every node $v \in L$ of the Label Cover instance, we have a set of k^n vertices denoted by $f_v = \{v\} \times [k]^n$. The vertex set of the hypergraph is $V' = \bigcup_{v \in L} f_v$.

For every $u \in R$, and k distinct neighbors of u , $v_1, v_2, \dots, v_k \in L$ with projection constraints $\pi_i : [\Sigma_L] \rightarrow [\Sigma_R], i \in [k]$, consider the set of k^2 vectors $\mathbf{x}^{i,j}$ for $i \in [k], j \in [k]$ which satisfy the following: For every $\beta \in \Sigma_R$, and for all $\alpha_1, \alpha_2, \dots, \alpha_k \in \Sigma_L$ such that $\pi_i(\alpha_i) = \beta$ for all $i \in [k]$, we have

$$|\{(i, j) | \mathbf{x}_{\alpha_i}^{i,j} = p\}| \leq 2k \quad \forall p \in [k] \quad (1)$$

For every such set of k^2 vectors, we add the edge $\{(v_i, \mathbf{x}^{i,j}) : 1 \leq i, j \leq k\}$ to E' . We can observe that $|V'| \leq |L|k^{|\Sigma_L|}$ and

$$|E'| \leq |R| \binom{\Delta}{k} \binom{k^{|\Sigma_L|}}{k}^k \leq |R| \Delta^k k^{|\Sigma_L|k^2}.$$

Completeness. Suppose that there exists an assignment $\sigma : V \rightarrow \Sigma$ that satisfies all the constraints of the Label Cover instance G . We color the set of vertices f_v in the long code corresponding to the vertex $v \in L$ with the dictator function on the coordinate $\sigma(v)$ i.e. for every $\mathbf{x} \in f_v$, we assign the color

$$c(\{v, \mathbf{x}\}) = \mathbf{x}_{\sigma(v)}$$

We can observe that this coloring satisfies the property that in every edge $e \in E'$, each color appears at most $2k$ times.

Soundness. Suppose that there is a proper k -coloring $c : V' \rightarrow [k]$ of the hypergraph H i.e. in every edge $e = \{v_1, v_2, \dots, v_{k^2}\}$, we have

$$|\{c(v_1), c(v_2), \dots, c(v_{k^2})\}| > 1$$

Our goal is to prove that there is a labeling to the Label Cover instance that satisfies at least $\epsilon = \frac{1}{k^8}$ fraction of constraints.

We need the following lemma proved by Austrin, Bhangale, Potukuchi [ABP20] using a generalization of Borsuk-Ulam theorem.

Lemma 4.20. (Theorem 5.2 of [ABP20]) *For every odd prime k and $n \geq k^3$, in any k -coloring of $[k]^n$, $c : [k]^n \rightarrow [k]$, there is a set of k vectors $\mathbf{x}^1, \mathbf{x}^2, \dots, \mathbf{x}^k$ that are all assigned the same color such that*

$$\{\mathbf{x}_i^1, \mathbf{x}_i^2, \dots, \mathbf{x}_i^k\} = [k]$$

for at least $n - k^3$ distinct coordinates $i \in [n]$.

Using this lemma, for every $v \in L$, we can identify a set of vectors $\mathbf{x}^{v,1}, \mathbf{x}^{v,2}, \dots, \mathbf{x}^{v,k} \in f_v$ such that all these vectors have the same color i.e. $c(\{v, \mathbf{x}^{v,i}\}) = c'(v)$ for all $v \in L, i \in [k]$ for some function $c' : L \rightarrow [k]$. Furthermore, there are a set of coordinates $S(v) \subseteq [n]$ with $|S(v)| \leq k^3$ such that

$$\{\mathbf{x}_i^{v,1}, \mathbf{x}_i^{v,2}, \dots, \mathbf{x}_i^{v,k}\} = [k]$$

for every $i \in [n] \setminus S(v)$.

For a set $S \subseteq \Sigma_L$ and a function $\pi : \Sigma_L \rightarrow \Sigma_R$, we use $\pi(S)$ to denote the set $\{\pi(i) : i \in S\}$. We now prove a key lemma that helps in the decoding procedure.

Lemma 4.21. *Let $u \in R$ be a node on the right side of the Label Cover instance. There are a set of labels $S(u) \subseteq \Sigma_R$ such that $|S(u)| \leq k^5$, and for every $v \in L$ that is a neighbor of u with projection constraint $\pi : \Sigma_L \rightarrow \Sigma_R$, we have $S(u) \cap \pi(S(v)) \neq \emptyset$.*

Proof. Fix a node $u \in R$ on the right side of the Label Cover instance. Let $v_1, v_2, \dots, v_l \in L$ be the neighbors of u in the Label Cover instance corresponding to the projection constraints $\pi_1, \pi_2, \dots, \pi_l$ respectively. As $|S(v_i)| \leq k^3$ for all $i \in [l]$, and the constraints π_i are projections, we have $|\pi_i(S(v_i))| \leq k^3$ for all $i \in [l]$. Among these l subsets $\pi_i(S(v_i))$ of Σ_R , let the maximum number of pairwise disjoint subsets be denoted by l' . Without loss of generality, we can assume that $\mathcal{S} = \{\pi_i(S(v_i)) : i \in [l']\}$ is a pairwise disjoint family of subsets.

We define the set $S(u)$ as follows:

$$S(u) = \bigcup_{i \in [l']} \pi_i(S(v_i))$$

As \mathcal{S} is a family of maximum pairwise disjoint subsets, we have $S(u) \cap \pi_i(S(v_i)) \neq \emptyset$ for all $i \in [l]$. Our goal is to bound the size of $S(u)$, which we achieve by bounding l' .

We claim that $l' \leq k(k-1)$. Suppose for contradiction that $l' > k(k-1)$. This implies that there are $l' > k(k-1)$ nodes $v_1, v_2, \dots, v_{l'}$ all adjacent to u such that $\pi_i(S(v_i)), i \in [l']$ are all pairwise disjoint. Thus, there exists a color $\ell \in [k]$ and a set of k nodes w_1, w_2, \dots, w_k adjacent to u corresponding to the projection constraints $\pi'_1, \pi'_2, \dots, \pi'_k$ such that $c'(w_i) = \ell$ for all $i \in [k]$, and the k sets $\pi'_i(S(w_i))$ are pairwise disjoint.

Using this, we can construct a set of vectors $\mathbf{x}^{i,j}, 1 \leq i, j \leq k$ defined as $\mathbf{x}^{i,j} = \mathbf{x}^{w_i,j}$ which satisfy the following properties:

1. All these vectors are colored the same:

$$c(\{w_i, \mathbf{x}^{i,j}\}) = \ell \quad \forall 1 \leq i, j \leq k$$

2. For every $i \in [k]$,

$$\{\mathbf{x}_{i'}^{i,1}, \mathbf{x}_{i'}^{i,2}, \dots, \mathbf{x}_{i'}^{i,k}\} = [k]$$

for every $i' \in [n] \setminus S(w_i)$.

We claim that these set of vectors satisfy the condition in Equation (1). Fix a $\beta \in \Sigma_R$, and $\alpha_1, \alpha_2, \dots, \alpha_k \in \Sigma_L$ such that $\pi'_i(\alpha_i) = \beta$ for all $i \in [k]$. As the family of subsets $\pi'_i(S(w_i))$ is a pairwise disjoint family, we can infer that there exists at most one $i \in [k]$ such that $\alpha_i \in S(w_i)$. Note that if $\alpha_i \notin S(w_i)$, then

$$\{\mathbf{x}_{\alpha_i}^{i,j} : j \in [k]\} = [k].$$

Thus, we have

$$|\{(i, j) | \mathbf{x}_{\alpha_i}^{i,j} = p\}| \leq 2k \quad \forall p \in [k].$$

Thus, the set of vectors $\{(w_i, \mathbf{x}^{i,j}) : 1 \leq i, j \leq k\}$ is indeed an edge of E' . As all these vectors are colored the same color ℓ , we have arrived at a contradiction to the fact that c is a proper k -coloring of H .

Hence, we can conclude that $l' \leq k(k-1)$, and thus, $|S(u)| \leq k(k-1)k^3 < k^5$. \square

Now, consider the labeling $\sigma : L \rightarrow \Sigma_L$, where $\sigma(v), v \in L$ is chosen uniformly at random from $S(v)$. Similarly, let $\sigma : R \rightarrow \Sigma_R$ is chosen uniformly at random from $S(u), u \in R$. Using Lemma 4.21, we can infer that for every edge $e = (v, u)$ in the Label Cover, this labeling satisfies the edge e with probability at least $\frac{1}{|S(v)||S(u)|} \geq \frac{1}{k^8}$. By linearity of expectation, this labeling satisfies at least $\frac{1}{k^8}$ fraction of the constraints in expectation. Hence, with positive probability, the labeling satisfies at least $\frac{1}{k^8}$ fraction of the constraints. This concludes the proof of soundness that if H has a proper k coloring, then there exists a labeling to G that satisfies at least $\frac{1}{k^8}$ fraction of the constraints.

5 Vector Bin Covering

5.1 Hardness of Vector Bin Covering via Rainbow Coloring

As is the case with the Vector Scheduling problem, the hard instances for Vector Bin Covering are when the vectors are from $\{0, 1\}^d$. In this setting, the Vector Bin Covering problem is closely related to the hypergraph rainbow coloring problem. A hypergraph $H = (V, E)$ is said to be k -rainbow colorable if there is an assignment of k colors to the vertices of H such that in every edge, all the k colors appear. When $k = 2$, it is equivalent to the standard 2-coloring of hypergraphs. Unlike the usual (hyper)graph coloring, rainbow coloring gets harder with larger number of colors.

In the approximate rainbow coloring problem, given a hypergraph that is promised to have a rainbow coloring with a large number of colors, the goal is to find a coloring in polynomial time using fewer number of colors. More formally, the computational problem we study is the following.

Definition 5.1. (*Approximate Rainbow Coloring*) *In the approximate rainbow coloring problem, the input is a hypergraph $H = (V, E)$ and a parameter $k \geq 2$. The objective is to distinguish between the following:*

1. *The hypergraph H is k -rainbow colorable.*

2. The hypergraph H has no valid 2-coloring.

We now give a simple reduction from approximate rainbow coloring to Vector Bin Covering.

Lemma 5.2. *Given a hypergraph $H = (V, E)$ and a parameter k , there is a polynomial time reduction that outputs a Vector Bin Covering instance $v_1, v_2, \dots, v_n \in \{0, 1\}^d$ with $n = |V|, d = |E|$ such that*

1. (Completeness.) *If H is k -rainbow colorable, there is a partition of $[n]$ into k parts A_1, A_2, \dots, A_k such that*

$$\sum_{j \in A_i} v_j \geq \mathbf{1}^d \quad \forall i \in [k]$$

2. (Soundness.) *If H is not 2-colorable, there is no partition of $[n]$ into A_1, A_2 such that*

$$\sum_{j \in A_i} v_j \geq \mathbf{1}^d \quad \forall i \in [2]$$

Proof. Let $n = |V|, d = |E|$. We order the edges E as $E = \{e_1, e_2, \dots, e_d\}$. We output a set of vectors $\mathcal{V} = \{v_1, v_2, \dots, v_n\}$ where each $v_i \in \{0, 1\}^d$ is defined as follows:

$$(v_i)_j = \begin{cases} 1 & \text{if } i \in e_j \\ 0 & \text{otherwise.} \end{cases}$$

We analyze this reduction:

1. (Completeness.) Suppose that the hypergraph H has a rainbow coloring with k colors $f : V \rightarrow [k]$. We partition $[n]$ into k parts A_1, A_2, \dots, A_k such that

$$A_i = \{j \in [n] : f(j) = i\}$$

Consider an arbitrary integer $i \in [k]$. Note that for every edge e in H , $e \cap A_i \neq \emptyset$. Thus,

$$\sum_{j \in A_i} v_j \geq \mathbf{1}^d$$

2. (Soundness.) Suppose that the hypergraph H has no proper coloring with 2 colors. Then, we claim that there is no partition of $[n]$ into two parts A_1, A_2 such that

$$\sum_{j \in A_i} v_j \geq \mathbf{1}^d \quad \forall i \in [2]$$

Suppose for contradiction that there exists A_1, A_2 with the above property. Consider the coloring of the hypergraph $f : V \rightarrow [2]$ as

$$f(v) = \begin{cases} 1 & \text{if } v \in A_1 \\ 2 & \text{if } v \in A_2 \end{cases}$$

Consider an arbitrary edge $e_l, l \in [d]$ of the hypergraph H . As $\sum_{j \in A_i} (v_j)_l \geq 1$ for all $i \in [2]$, there exist $v_1, v_2 \in e_l$ such that $v_1 \in A_1, v_2 \in A_2$. Thus, the coloring f is a proper 2 coloring of the hypergraph H , a contradiction. \square

We combine this reduction with the hardness of approximate rainbow coloring to prove the hardness of Vector Bin Covering, namely Theorem 1.5. Note that the dimension of the resulting vectors in the Vector Bin Covering instance \mathcal{V} is equal to the number of edges $m = |E|$ of the hypergraph H , and the gap in the optimal Bin Covering value of \mathcal{V} is equal to k , the number of colors. Hence, to obtain better inapproximability results for Vector Bin Covering that grow with d , our goal is to show the hardness of approximate rainbow coloring on hypergraphs with m edges where the number of colors k is as large a function of m as possible. Towards this, we prove that it is NP-hard to 2-color a hypergraph with m edges that is promised to be rainbow colorable with $k = \Omega\left(\frac{\log m}{\log \log m}\right)$ colors.

Theorem 5.3. *Given a hypergraph H with m edges, it is NP-hard to distinguish between the following:*

1. (Completeness) H is k -rainbow colorable.
2. (Soundness) H is not 2-colorable.

where $k = \Omega\left(\frac{\log m}{\log \log m}\right)$.

We defer the proof of Theorem 5.3 to Section 5.2.

We now prove the hardness of Vector Bin Covering using Theorem 5.3.

Proof of Theorem 1.5. Using Theorem 5.3 combined with the reduction in Lemma 5.2, we get that the following problem is NP-hard. Given a set of n vectors $v_1, v_2, \dots, v_n \in \{0, 1\}^d$, distinguish between

1. \mathcal{V} can be partitioned into $k = \Omega\left(\frac{\log d}{\log \log d}\right)$ parts such that in each part, the sum of vectors is at least 1 in every coordinate.
2. \mathcal{V} cannot be partitioned into 2 parts such that in each part, the sum of vectors is at least 1 in every coordinate. In other words, the maximum number of parts into which \mathcal{V} can be partitioned such that in each part, the sum of vectors is at least 1 in every coordinate is equal to 1.

Thus, it is NP-hard to approximate d -dimensional Vector Bin Covering within $k = \Omega\left(\frac{\log d}{\log \log d}\right)$. \square

5.2 Proof of Theorem 5.3

Our proof follows by viewing the hypergraph rainbow coloring problem as a promise constraint satisfaction problem (PCSP) and analyzing its polymorphisms [AGH17, BG16, BKO19]. The idea is to prove that the polymorphisms have a small number of “important” coordinates which can then be decoded in the Label Cover-Long Code framework. For the case of the above rainbow coloring PCSP, we prove that the polymorphisms are 1-fixing in that there is a single coordinate which when set to a certain value fixes the value of the function. This characterization then implies the hardness of the approximate rainbow coloring.

For ease of readability, we skip defining polymorphisms formally, and instead present the proof as a simple gadget reduction from Label Cover. We first need a definition.

Definition 5.4. (1-fixing [BG16, GS20]) *A function $f : [k]^n \rightarrow \{0, 1\}$ is said to be 1-fixing if there exists an index $\ell \in [n]$ and values $\alpha, \beta \in [k]$ such that*

$$f(\mathbf{x}) = 0 \forall \mathbf{x} : \mathbf{x}_\ell = \alpha \quad \text{and} \quad f(\mathbf{x}) = 1 \forall \mathbf{x} : \mathbf{x}_\ell = \beta$$

In the analysis of the gadget, we need a definition and a lemma from [ABP20].

Definition 5.5. (The hypergraph $H_r^n[k]$) The hypergraph $H_r^n[k] = (V, E)$ is a k -uniform hypergraph with vertex set as the set of n -dimensional vectors over $[k]$ i.e. $V = [k]^n$. A set of k vectors $\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^k$ form an edge of the hypergraph if

$$\sum_{i=1}^n \left| [k] \setminus \{\mathbf{v}_i^j : j \in [k]\} \right| \leq r$$

Lemma 5.6. For every $k \geq 2$, the hypergraph $H_{\lfloor \frac{k}{2} \rfloor}^n[k]$ is not 2-colorable.

We analyze the gadget used in our reduction.

Lemma 5.7. Fix $k \geq 3$. Suppose $f : [k]^n \rightarrow \{0, 1\}$ satisfies the below two-coloring property: For every $2k$ vectors $\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^{2k} \in [k]^n$ with

$$\{\mathbf{v}_i^j : j \in [2k]\} = [k] \forall i \in [n],$$

we have

$$\{f(\mathbf{v}^j) : j \in [2k]\} = \{0, 1\}.$$

Then, f is 1-fixing.

Proof. We first prove that there exist $\ell \in [n], \alpha \in [k], b \in \{0, 1\}$ such that $f(\mathbf{x}) = b$ for all $\mathbf{x} \in [k]^n$ with $\mathbf{x}_\ell = \alpha$. Suppose for contradiction that this is not the case. Then, for every $i \in [n], j \in [k]$ there exist vectors $\mathbf{x}^{i,j}, \mathbf{y}^{i,j} \in [k]^n$ such that $\mathbf{x}_i^{i,j} = \mathbf{y}_i^{i,j} = j$, and $f(\mathbf{x}^{i,j}) = 0$ where as $f(\mathbf{y}^{i,j}) = 1$.

Let $r = \lfloor \frac{k}{2} \rfloor$. We view $f : [k]^n \rightarrow \{0, 1\}$ as an assignment of two colors to the vertices of the hypergraph $H_r^n[k]$. As the hypergraph is not two colorable (Lemma 5.6), we can infer that there is an edge of $H_r^n[k]$ all of whose vertices are assigned the same color. In other words, there exist k vectors $\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^k \in [k]^n$ and $b \in \{0, 1\}$ such that $f(\mathbf{v}^j) = b$ for all $j \in [k]$. Furthermore, there are at most r missing values in these vectors i.e.

$$\sum_{i=1}^n \left| [k] \setminus \{\mathbf{v}_i^j : j \in [k]\} \right| \leq r$$

Now, we pick r vectors $\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^r$ (with repetitions if needed) by filling the missing values using $\mathbf{x}^{i,j}, \mathbf{y}^{i,j}$ vectors such that

1. $f(\mathbf{u}^j) = b$ for all $j \in [r]$.

2. For every $i \in [n]$,

$$\{\mathbf{v}_i^j : j \in [k]\} \cup \{\mathbf{u}_i^j : j \in [r]\} = [k]$$

By taking the union of $\{\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^k\}$ and $\{\mathbf{u}^1, \mathbf{u}^2, \dots, \mathbf{u}^r\}$, and repeating some vectors, we obtain $2k$ vectors $\mathbf{w}^1, \mathbf{w}^2, \dots, \mathbf{w}^{2k}$ with $f(\mathbf{w}^j) = b$ for all $j \in [2k]$, and

$$\{\mathbf{w}_i^j : j \in [2k]\} = [k] \forall i \in [n]$$

However, this contradicts the two-coloring property of f . Thus, there exist $\ell \in [n], \alpha \in [k], b \in \{0, 1\}$ such that $f(\mathbf{x}) = b$ for all $\mathbf{x} \in [k]^n$ with $\mathbf{x}_\ell = \alpha$.

We now claim that there exists $\beta \in [k]$ such that $f(\mathbf{x}) = 1 - b$ for all $\mathbf{x} \in [k]^n$ with $\mathbf{x}_\ell = \beta$. Suppose for contradiction that this is not the case. Then, there exist k vectors $\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^k$ such that $\mathbf{v}_\ell^j = j$ for all $j \in [k]$, and $f(\mathbf{v}^j) = b$ for all $j \in [k]$. We now pick $\mathbf{v}^{k+1}, \mathbf{v}^{k+2}, \dots, \mathbf{v}^{2k} \in [k]^n$ such that $\mathbf{v}_\ell^j = \alpha$ for all $j \in \{k+1, k+2, \dots, 2k\}$, and $\mathbf{v}_i^j = j - k$ for all $i \in [n]$ with $i \neq \ell$, and $j \in \{k+1, k+2, \dots, 2k\}$. These $2k$ vectors $\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^{2k}$ satisfy

1. $f(\mathbf{v}^j) = b$ for all $j \in [2k]$.
2. For every $i \in [n]$,

$$\{\mathbf{v}_i^j : j \in [2k]\} = [k]$$

contradicting the two-coloring property of f . Thus, there exists $\beta \in [k]$ such that $f(\mathbf{x}) = 1 - b$ for all $\mathbf{x} \in [k]^n$ with $\mathbf{x}_\ell = \beta$, completing the proof that f is 1-fixing. \square

We are now ready to prove Theorem 5.3. Our hardness result is obtained using a reduction from the Label Cover problem. This reduction is standard in the PCSP literature. (See e.g., [BKO19])

Reduction. We start with the Label Cover instance $G = (V = L \cup R), E, \Sigma = \Sigma_L = \Sigma_R, \Pi)$ from Theorem 2.6 and output a hypergraph $H = (V', E')$. Let n denote the label size $n = |\Sigma|$. For each vertex $v \in L \cup R$, we have a long code containing a set of nodes K_v of size $[k]^n$, indexed by n length vectors.

1. The vertex set of the hypergraph V' is the union of all the long code nodes.

$$V' = \bigcup_{v \in V} K_v$$

2. Edges of the hypergraph: For every vertex $v \in V$ of the Label Cover instance, we add an edge in E' for each set of $2k$ vectors $\{\mathbf{v}^1, \mathbf{v}^2, \dots, \mathbf{v}^{2k}\}$ in K_v , if

$$\{\mathbf{v}_i^j : j \in [2k]\} = [k] \forall i \in [n]. \quad (2)$$

The number of edges in H is at most

$$|E'| \leq |V| \binom{k^{|\Sigma|}}{2k} \leq |V| k^{O(k)}$$

3. Equality constraints: For every constraint $\Pi_e : u \rightarrow v$ of the Label Cover, we add a set of equality constraints between nodes $\mathbf{x} \in K_u, \mathbf{y} \in K_v$ if for all $i \in [n]$, $\mathbf{x}_i = \mathbf{y}_{\Pi_e(i)}$. By adding an equality constraint between two nodes, we identify the two nodes together and treat it as a single node. That is, we compute the connected components of the equality constraints graph and identify a single master node for each component. We then obtain a multi-hypergraph H_1 from H by replacing each node with the corresponding master node. However, a vertex could appear multiple times in an edge in H_1 . We delete such occurrences from H_1 by setting each edge to be a simple set of the vertices contained in it, and obtain the final hypergraph H_2 . We note the following:

- (a) There exists a k -rainbow coloring of H , $f : V' \rightarrow [k]$ that respects the equality constraints i.e. $f(\mathbf{x}) = f(\mathbf{y})$ for all pairs of nodes \mathbf{x}, \mathbf{y} with equality constraints between them if and only if H_2 is k -rainbow colorable.
- (b) Similarly, there exists a 2-coloring of H that respects equality constraints if and only if H_2 is 2-colorable.

Finally, the number of edges in H_2 is at most the number of edges in H .

Completeness. Suppose that there is a labeling $\sigma : V \rightarrow \Sigma$ that satisfies all the constraints. We define the coloring $f : V' \rightarrow [k]$ of H as follows. For every node $\mathbf{x} \in K_v$, we set the dictatorship function

$$f(\mathbf{x}) = \mathbf{x}_{\sigma(v)}$$

By the constraints added in Equation (2), the function f is a valid k -rainbow coloring of H . As σ satisfies all the constraints of the Label Cover, the coloring f satisfies all the equality constraints.

Soundness. Suppose that there is no labeling $\sigma : V \rightarrow \Sigma$ that satisfies all the constraints in G . Then we claim that there is no 2-coloring of H that respects all the equality constraints. Suppose for contradiction that there is a 2-coloring $f : V' \rightarrow \{0, 1\}$ that respects all the equality constraints.

Consider a vertex $v \in V$. The function $f_v : [k]^n \rightarrow \{0, 1\}$, defined as f on K_v satisfies the conditions in Lemma 5.7. Thus, f_v is 1-fixing for every $v \in V$. Hence, there is a function $L : V \rightarrow \Sigma$ such that for every $v \in V$, f_v is 1-fixing on the coordinate $L(v)$. We now claim that the labeling $\sigma : V \rightarrow \Sigma$ defined as $\sigma(v) = L(v)$ satisfies all the constraints in G .

Consider an edge $e = (u, v)$, $u \in L, v \in R$ with the projection constraint $\Pi_e : \Sigma \rightarrow \Sigma$. Our goal is to show that $\Pi_e(L(u)) = L(v)$. Suppose for contradiction that $\Pi_e(L(u)) \neq L(v)$. By the 1-fixing property of f_u , we have $\alpha_u, \beta_u \in [k]$ such that

$$f_u(\mathbf{x}) = 0 \forall \mathbf{x} \in [k]^n : \mathbf{x}_{L(u)} = \alpha_u \quad \text{and} \quad f_u(\mathbf{x}) = 1 \forall \mathbf{x} \in [k]^n : \mathbf{x}_{L(u)} = \beta_u$$

Similarly, we have $\alpha_v, \beta_v \in [k]$ such that

$$f_v(\mathbf{y}) = 0 \forall \mathbf{y} \in [k]^n : \mathbf{y}_{L(v)} = \alpha_v \quad \text{and} \quad f_v(\mathbf{y}) = 1 \forall \mathbf{y} \in [k]^n : \mathbf{y}_{L(v)} = \beta_v$$

By the equality constraints, $f_u(\mathbf{x}) = f_v(\mathbf{y})$ for all $\mathbf{x}, \mathbf{y} \in [k]^n$ such that $\mathbf{x}_i = \mathbf{y}_{\Pi_e(i)} \forall i \in [n]$. Let $\mathbf{y}' \in [k]^n$ be an arbitrary vector with $\mathbf{y}'_{\Pi_e(L(u))} = \alpha_u, \mathbf{y}'_{L(v)} = \beta_v$. We choose $\mathbf{x}' \in [k]^n$ such that for all $i \in [n]$, $\mathbf{x}'_i = \mathbf{y}'_{\Pi_e(i)}$. Note that $\mathbf{x}'_{L(u)} = \alpha_u$. Thus, $f_u(\mathbf{x}') = 0$ where as $f_v(\mathbf{y}') = 1$. However, this contradicts the equality constraints.

6 Conclusion

We conclude by mentioning a few open problems.

1. A drawback of our hardness result for Vector Bin Packing is that our lower bound is only applicable when d is large enough. In particular, for the 2-dimensional Vector Bin Packing, we still do not have an explicit constant hardness. The reason our hardness result needs d to be large enough is that the k -set cover hardness itself [Tre01] needs k to be large enough. By starting our reduction with alternate set cover hardness results such as the 3-dimensional matching problem, we can obtain improved hardness for smaller values of d . However, this approach will still not help for $d = 2$ as the upper bound on the packing dimension of these set families is greater than 2. It is an interesting open problem to characterize set families with packing dimension 2. We believe that such a characterization could help in proving the hardness of approximation results for the set cover problem on set families with packing dimension 2, which directly gives the hardness of 2-dimensional Vector Bin Packing.
2. d -dimensional Geometric Bin Packing is another open problem where packing dimension based ideas could help. The best hardness result for the d -dimensional Geometric Bin Packing is still the (tiny) hardness factor from the 2-dimensional setting. A possible avenue to obtain improved inapproximability for this problem is by a reduction from a suitable set cover variant using a notion of packing dimension. However, this is easier said than done as the geometric packings are significantly harder to tame—for example, it is NP-hard to decide if a given set of n rectangles can fit in a unit square, while the corresponding problem for Vector Bin Packing is trivial.

3. Another interesting open problem is to close the gap between $O(\log d)$ algorithm and $\Omega\left(\frac{\log d}{\log \log d}\right)$ hardness for Vector Bin Covering. For the rainbow coloring, the question is: Given a hypergraph H with m edges that is promised to be $f(m)$ -rainbow colorable, can we 2-color it in polynomial time? The answer to this question when $f(m)$ is equal to $O(\log m)$ is yes, by a simple random 2-coloring. By Theorem 5.3, the answer is no when $f(m) = \Omega\left(\frac{\log m}{\log \log m}\right)$. Which of these is tight? Bhattachiprolu, Guruswami, and Lee [BGL15] proved that in certain settings, the simple random coloring is optimal for rainbow coloring. This suggests that perhaps even here, the problem is hard when $f(m) = o(\log m)$. On the other hand, obtaining any hardness result beyond $\Omega\left(\frac{\log m}{\log \log m}\right)$ seems to be impossible with the Label Cover-Long Code framework.

Acknowledgements

I am greatly indebted to Venkatesan Guruswami for helpful discussions, for his detailed feedback on the manuscript which significantly improved the presentation, and for his encouragement. I also thank Varun Gupta, Ravishankar Krishnaswamy, and Janani Sundaresan for helpful discussions. I am especially grateful to Ravishankar Krishnaswamy for pointing me to [KAR00].

References

- [AAC⁺98] Noga Alon, Yossi Azar, János Csirik, Leah Epstein, Sergey V. Sevastianov, Arjen P. A. Vestjens, and Gerhard J. Woeginger. On-line and off-line approximation algorithms for vector covering problems. *Algorithmica*, 21(1):104–118, 1998.
- [ABP19] Per Austrin, Amey Bhangale, and Aditya Potukuchi. Simplified inapproximability of hypergraph coloring via t-agreeing families. *CoRR*, abs/1904.01163, 2019.
- [ABP20] Per Austrin, Amey Bhangale, and Aditya Potukuchi. Improved inapproximability of rainbow coloring. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020*, pages 1479–1495, 2020.
- [ACKS13] Yossi Azar, Ilan Reuven Cohen, Seny Kamara, and F. Bruce Shepherd. Tight bounds for online vector bin packing. In *Symposium on Theory of Computing Conference, STOC’13*, pages 961–970, 2013.
- [AGH17] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\epsilon)$ -sat is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017.
- [AJKL84] Susan F. Assmann, David S. Johnson, Daniel J. Kleitman, and Joseph Y.-T. Leung. On a dual version of the one-dimensional bin packing problem. *J. Algorithms*, 5(4):502–525, 1984.
- [ALM⁺98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. Proof verification and the hardness of approximation problems. *J. ACM*, 45(3):501–555, 1998.
- [Ban17] Nikhil Bansal. Scheduling open problems: Old and new. *MAPSP 2017*, www.mapsp2017.ma.tum.de/MAPSP2017-Bansal.pdf, 2017.

- [BCS09] Nikhil Bansal, Alberto Caprara, and Maxim Sviridenko. A new approximation method for set covering problems, with applications to multidimensional bin packing. *SIAM J. Comput.*, 39(4):1256–1278, 2009.
- [BEK16] Nikhil Bansal, Marek Eliáš, and Arindam Khan. Improved approximation for vector bin packing. In *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pages 1561–1579, 2016.
- [BG16] Joshua Brakensiek and Venkatesan Guruswami. New hardness results for graph and hypergraph colorings. In *31st Conference on Computational Complexity, CCC 2016*, pages 14:1–14:27, 2016.
- [BGL15] Vijay V. S. P. Bhattiprolu, Venkatesan Guruswami, and Euiwoong Lee. Approximate hypergraph coloring under low-discrepancy and related promises. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2015*, pages 152–174, 2015.
- [Bha18] Amey Bhangale. NP-Hardness of coloring 2-colorable hypergraph with polylogarithmically many colors. In *45th International Colloquium on Automata, Languages, and Programming, ICALP 2018*, pages 15:1–15:11, 2018.
- [BK14] Nikhil Bansal and Arindam Khan. Improved approximation algorithm for two-dimensional bin packing. In *Proceedings of the Twenty-Fifth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014*, pages 13–25, 2014.
- [BKO19] Jakub Bulín, Andrei A. Krokhin, and Jakub Oprsál. Algebraic approach to promise constraint satisfaction. In *Proceedings of the 51st Annual ACM SIGACT Symposium on Theory of Computing, STOC 2019*, pages 602–613, 2019.
- [BOVvdZ16] Nikhil Bansal, Tim Oosterwijk, Tjark Vredeveld, and Ruben van der Zwaan. Approximating vector scheduling: Almost matching upper and lower bounds. *Algorithmica*, 76(4):1077–1096, 2016.
- [BS04] Nikhil Bansal and Maxim Sviridenko. New approximability and inapproximability results for 2-dimensional bin packing. In *Proceedings of the Fifteenth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2004*, pages 196–203, 2004.
- [CC06] Miroslav Chlebík and Janka Chlebíková. Inapproximability results for orthogonal rectangle packing problems with rotations. In *Algorithms and Complexity, 6th Italian Conference, CIAC 2006*, pages 199–210, 2006.
- [CFS08] L. Sunil Chandran, Mathew C. Francis, and Naveen Sivadasan. Boxicity and maximum degree. *J. Comb. Theory, Ser. B*, 98(2):443–445, 2008.
- [CJK01] János Csirik, David S. Johnson, and Claire Kenyon. Better approximation algorithms for bin covering. In *Proceedings of the Twelfth Annual Symposium on Discrete Algorithms, 2001*, pages 557–566, 2001.
- [CK04] Chandra Chekuri and Sanjeev Khanna. On multidimensional packing problems. *SIAM J. Comput.*, 33(4):837–851, 2004.
- [CKPT17] Henrik I. Christensen, Arindam Khan, Sebastian Pokutta, and Prasad Tetali. Approximation and online algorithms for multidimensional bin packing: A survey. *Comput. Sci. Rev.*, 24:63–79, 2017.

- [dIVL81] Wenceslas Fernandez de la Vega and George S. Lueker. Bin packing can be solved within $1+\epsilon$ in linear time. *Comb.*, 1(4):349–355, 1981.
- [DS14] Irit Dinur and David Steurer. Analytical approach to parallel repetition. In *Symposium on Theory of Computing, STOC 2014*, pages 624–633, 2014.
- [Fei98] Uriel Feige. A threshold of $\ln n$ for approximating set cover. *J. ACM*, 45(4):634–652, 1998.
- [FK15] Alan Frieze and Michał Karoński. *Introduction to Random Graphs*. Cambridge University Press, 2015.
- [FKT89] Ulrich Faigle, Walter Kern, and György Turán. On the performance of on-line algorithms for partition problems. *Acta Cybern.*, 9(2):107–119, 1989.
- [GGJY76] M. R. Garey, Ronald L. Graham, David S. Johnson, and Andrew Chi-Chih Yao. Resource constrained scheduling as generalized bin packing. *J. Comb. Theory, Ser. A*, 21(3):257–298, 1976.
- [GJ79] M. R. Garey and David S. Johnson. *Computers and Intractability: A Guide to the Theory of NP-Completeness*. W. H. Freeman, 1979.
- [GL18] Venkatesan Guruswami and Euiwoong Lee. Strong inapproximability results on balanced rainbow-colorable hypergraphs. *Comb.*, 38(3):547–599, 2018.
- [GS20] Venkatesan Guruswami and Sai Sandeep. Rainbow coloring hardness via low sensitivity polymorphisms. *SIAM J. Discret. Math.*, 34(1):520–537, 2020.
- [HS87] Dorit S. Hochbaum and David B. Shmoys. Using dual approximation algorithms for scheduling problems theoretical and practical results. *J. ACM*, 34(1):144–162, 1987.
- [HS19] David G. Harris and Aravind Srinivasan. The Moser-Tardos framework with partial resampling. *J. ACM*, 66(5):36:1–36:45, 2019.
- [IKKP19] Sungjin Im, Nathaniel Kell, Janardhan Kulkarni, and Debmalya Panigrahi. Tight bounds for online vector scheduling. *SIAM J. Comput.*, 48(1):93–121, 2019.
- [Jan10] Klaus Jansen. An EPTAS for scheduling jobs on uniform processors: Using an MILP relaxation with a constant number of integral variables. *SIAM J. Discret. Math.*, 24(2):457–485, 2010.
- [JS03] Klaus Jansen and Roberto Solis-Oba. An asymptotic fully polynomial time approximation scheme for bin covering. *Theor. Comput. Sci.*, 306(1-3):543–551, 2003.
- [Juh82] Ferenc Juhász. The asymptotic behaviour of Lovász’ theta-function for random graphs. *Comb.*, 2(2):153–155, 1982.
- [KAR00] V. S. Anil Kumar, Sunil Arya, and H. Ramesh. Hardness of set cover with intersection 1. In *Automata, Languages and Programming, 27th International Colloquium, ICALP 2000*, pages 624–635, 2000.
- [Kho01] Subhash Khot. Improved inapproximability results for maxclique, chromatic number and approximate graph coloring. In *42nd Annual Symposium on Foundations of Computer Science, FOCS 2001*, pages 600–609, 2001.

- [Knu94] Donald E. Knuth. The sandwich theorem. *Electron. J. Comb.*, 1, 1994.
- [LY94] Carsten Lund and Mihalis Yannakakis. On the hardness of approximating minimization problems. *J. ACM*, 41(5):960–981, 1994.
- [Mos15] Dana Moshkovitz. The projection games conjecture and the NP-hardness of \ln n -approximating set-cover. *Theory Comput.*, 11:221–235, 2015.
- [MR10] Dana Moshkovitz and Ran Raz. Two-query PCP with subconstant error. *J. ACM*, 57(5):29:1–29:29, 2010.
- [MRT13] Adam Meyerson, Alan Roytman, and Brian Tagiku. Online multidimensional load balancing. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2013*, pages 287–302, 2013.
- [PTUW11] Rina Panigrahy, Kunal Talwar, Lincoln Uyeda, and Udi Wieder. Heuristics for vector bin packing. 2011.
- [Raz98] Ran Raz. A parallel repetition theorem. *SIAM J. Comput.*, 27(3):763–803, 1998.
- [Spi94] Frits CR Spieksma. A branch-and-bound algorithm for the two-dimensional vector packing problem. *Computers & operations research*, 21(1):19–25, 1994.
- [ST12] Hadas Shachnai and Tami Tamir. Approximation schemes for generalized two-dimensional vector packing with application to data placement. *J. Discrete Algorithms*, 10:35–48, 2012.
- [Tre01] Luca Trevisan. Non-approximability results for optimization problems on bounded degree instances. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, 2001*, pages 453–461, 2001.
- [Woe97] Gerhard J. Woeginger. There is no asymptotic PTAS for two-dimensional vector packing. *Inf. Process. Lett.*, 64(6):293–297, 1997.

A Hardness of simple k -set cover

The hardness result of Kumar, Arya, and Ramesh [KAR00] is obtained from the Label Cover problem using a partition gadget along the lines of the reduction of Lund and Yannakakis [LY94]. The set families in the reduction in [LY94] have large intersections. [KAR00] get around this by using two main ideas:

1. They use a different partition system wherein each partition is a disjoint union of a large (super constant) number of sets instead of just 2 sets in [LY94].
2. They use multiple sets for each label assignment to a vertex of the Label Cover, unlike a single set corresponding to each label of each vertex in [LY94].

As [KAR00] were proving a $\Omega(\log n)$ hardness of the set cover, the universe size of the partition system is chosen to be the same as the number of vertices in the Label Cover instance. This forces the set sizes to be very large. We can get around this issue by simply defining the partition system on a set of size B , where B is a large constant. This also has an added benefit that we no longer require sub-constant hardness from the Label Cover instances, thus giving us NP-hardness directly. This observation is used by

Trevisan [Tre01] to obtain $\ln B - O(\ln \ln B)$ NP-hardness of set cover on instances where each set has cardinality at most B , from Feige's $(1 - \epsilon) \ln n$ set cover hardness [Fei98].

We now describe the parameter modifications in full detail. Let B be a large constant.

We start our reduction from Label Cover instances with soundness $\gamma = \frac{1}{32\beta^2 \log^2 B}$ where β is an absolute constant to be fixed later.

Theorem A.1. ([ALM⁺98, Raz98]) *Given a Label Cover instance defined on a bipartite graph $G = (V, E)$ with left alphabet Σ_L and right alphabet Σ_R , it is NP-hard to distinguish between the following:*

1. (Completeness). *There exists a labeling $\sigma : V \rightarrow \Sigma_L \cup \Sigma_R$ that satisfies all the constraints.*
2. (Soundness). *No labeling to V can satisfy more than γ fraction of the constraints.*

Furthermore the instances satisfy the following properties:

1. *The alphabet sizes $d = |\Sigma_L|$ and $d' = |\Sigma_R|$ are both upper bounded by $(\log B)^{O(1)}$.*
2. *The maximum degree \deg of G is upper bounded by $(\log B)^{O(1)}$.*

Following the convention in [KAR00], we assume that the number of vertices on the left side in G is equal to that on the right side of G , and we denote this number by n' .

We now construct a partition system \mathcal{P} on a universe N of size B . The system \mathcal{P} has $d' \times (\deg + 1) \times d$ partitions. Each partition has $m = B^{1-\epsilon}$ parts, where ϵ is a small constant to be fixed later. The partition system is divided into d' groups each containing $(\deg + 1) \times d$ partitions. Each group is further organized into $\deg + 1$ subgroups each of which contains d partitions. Let $P_{g,s,p}$ denote the p th partition in the s th subgroup of the g th group and $P_{g,s,p,k}$ denote the k th set in $P_{g,s,p}$ where $g \in [d']$, $s \in [\deg + 1]$, $p \in [d]$, $k \in [m]$. The partition system satisfies the four properties in Section 4 of [KAR00], the only difference being that the universe N now has size B instead of n' . Thus, the covering property (Property 4 in [KAR00]) now states that any covering of N with $\beta m \log B$ sets should contain at least $\frac{3m}{4}$ sets from the same partition. Such a partition system is shown to exist for large enough B in [KAR00] using a randomized construction. They also derandomize the construction. But for our setting, as B is a constant, we just need to show the existence of such a partition system.

We reduce the Label Cover instance in Theorem A.1 to a set cover instance \mathcal{SC} by the same construction as in [KAR00]: we have a partition system corresponding to each edge of the Label Cover instance, and the union of the elements in the partition systems is the element set of \mathcal{SC} . The sets in \mathcal{SC} , $C_k(v, a)$ are defined exactly as in [KAR00]. The cardinality of each set is at most $B' = \deg \times B \leq B^2$. Each element is present in at most $md = O(B)$ sets. The fact that \mathcal{SC} is a simple set system follows from Lemma 1 of [KAR00]. By Lemma 2 in [KAR00], if there is a labeling of the Label Cover instance, then there is a set cover of size $n'm$ in \mathcal{SC} . If there is a set cover of size $\frac{\beta}{2} n'm \log B$ in \mathcal{SC} , then there is a labeling of G that satisfies γ fraction of constraints. The proof of this soundness follows along the same lines as Lemma 3 of [KAR00], with the only difference being that we now define the *good* edges as edges having $\#(e) \leq \beta m \log B$.

B SDP Relaxation of Monochromatic-Clique

We consider the following SDP relaxation of the graph coloring problem on $G = (V, E)$:

Minimize k

$$\langle u_i, u_i \rangle = 1 \quad \forall i \in V$$

$$\langle u_i, u_j \rangle \leq \frac{-1}{k-1} \quad \forall (i, j) \in E$$

The optimal solution to this SDP is referred to as the vector chromatic number $\chi_v(G)$ of the graph G . It is equivalent to the Lovasz theta function of the complement of G . We have the following sandwich property due to [Knu94]:

$$\omega(G) \leq \chi_v(G) \leq \chi(G)$$

B.1 Algorithm when $B > \sqrt{n}$

There is a simple algorithm for the Monochromatic-Clique(k, B) problem when $B > k$: We compute $\chi_v(G)$ in polynomial time, and we check if $\chi_v(G) \leq k$. In this case, there is no clique of size B in G , and we output YES. If $\chi_v(G) > k$, then the graph cannot be colored with k colors, and in this case, we output NO.

Note that if $k(B - 1) \geq n$, there is always an assignment of k colors to the vertices of the graph without a clique of size B , thus the problem is trivial.

B.2 Integrality gap

The above algorithm proves that in any graph with vector chromatic number at most k , there is an assignment of k colors to the vertices that has monochromatic clique of size at most \sqrt{n} . We now prove that this cannot be significantly improved:

Theorem B.1. *For n large enough, there exists a graph $G = (V, E)$ with n vertices, and a parameter k such that*

1. $\chi_v(G) \leq k$.
2. *In any assignment of k colors to the vertices of G , there is a monochromatic clique of size $n^{\Omega(1)}$.*

Proof. We first prove the following: for large enough n , there exists a graph G on n vertices, and an integer k such that

1. $\vartheta(G) \leq k$.
2. In any assignment of k colors to the vertices of the graph G , there exists a monochromatic independent set of size $B = n^{\Omega(1)}$.

Our construction is a probabilistic one: we sample G from $G(n, p)$ with $p = \frac{1}{\sqrt{n}}$. It has been proved [Juh82] that the Lovasz theta function of this random graph satisfies

$$\vartheta(G) \leq 2n^{\frac{3}{4}} + \tilde{O}(n^{\frac{1}{3}} \log n)$$

with high probability. We set $k = 3n^{\frac{3}{4}}$. For large enough n , with high probability, we have $\vartheta(G) \leq k$.

Furthermore, the random graph $G(n, p)$ with $p = o(n^{-\frac{2}{5}})$ has no K_6 with high probability (See e.g., [FK15]). Thus, using Lemma 4.2, we can infer that in any subset of size $\frac{n^{\frac{1}{3}}}{3}$, there is an independent set of size at least $\frac{n^{\frac{1}{24}}}{2}$. Hence, in any assignment of k colors to the vertices of the graph G , there is a monochromatic independent set of size $n^{\Omega(1)}$. Taking the complement, we get a graph with the required properties. \square