

New Directions in Inapproximability: Promise Constraint Satisfaction Problems and Beyond

Sai Sandeep Reddy Pallerla

CMU-CS-22-139

August 2022

School of Computer Science
Carnegie Mellon University
Pittsburgh, PA 15213

Thesis Committee:

Venkatesan Guruswami, Chair

Anupam Gupta

Ryan O'Donnell

Nikhil Bansal, University of Michigan

*Submitted in partial fulfillment of the requirements
for the degree of Doctor of Philosophy.*

Copyright © 2022 Sai Sandeep Reddy Pallerla

This research was sponsored by Google, The David and Lucille Packard Foundation under award number 200529094A, and the National Science Foundation under award numbers CCF-1422045, CCF-1563742, CCF-1814603, and CCF-1908125. The views and conclusions contained in this document are those of the author and should not be interpreted as representing the official policies, either expressed or implied, of any sponsoring institution, the U.S. government or any other entity.

Keywords: Approximation Algorithms, Hardness of Approximation, Approximate Graph Coloring, Promise Constraint Satisfaction Problems, Polymorphisms, Hypergraph Rainbow Coloring, Semi Definite Programming Relaxations, Set Cover, Vector Bin Packing, Hypergraph Vertex Cover

To my parents and my sister

Abstract

The field of hardness of approximation has seen a lot of progress in the past three decades resulting in almost optimal inapproximability results for many computational problems, including all Constraint Satisfaction Problems (CSPs). However, our understanding of inapproximability is still rather limited for some fundamental problems such as approximate graph coloring, and problems for which approximation algorithms have been widely studied, for example, clustering, packing, and scheduling problems. In this thesis, we make progress on these questions by studying Promise CSPs that generalize CSPs, and more abstractly, computational problems with a given promise, either a solution satisfying a strong property, or a structural guarantee on the underlying instance.

Promise Constraint Satisfaction Problems (PCSPs) generalize the traditional CSPs by allowing for a weaker and stronger form for each predicate. PCSPs have received a lot of attention recently, both on the algorithmic and hardness front, and their study has led to breakthroughs in approximate graph and hypergraph coloring problems. In this thesis, we continue that line of work, obtaining new characterization of polynomial time solvability for several classes of Promise CSPs including Boolean monotone PCSPs, variants of graph and hypergraph colorings. We also study robust algorithms for Promise CSPs, and give a dichotomy result characterizing when certain classes of PCSPs have robust algorithms. More generally, we study combinatorial problems under a given structural promise – for example, set cover on set systems where every pair of sets intersect in at most one element. We use inapproximability results on such structured instances to resolve the approximability of multidimensional packing problems and scheduling with communication delays.

Acknowledgements

First and foremost, I feel deeply fortunate to work with Venkat. He has been a wonderful advisor who adapts to the students' needs, mentoring them closely at times, and giving them freedom when needed. I am thankful to him for all the countless hours spent working on research, his superhuman patience, and for being an inspiration, both as a researcher and a person. Apart from the technical aspects, I have learned many things working with him, including how to pick problems to work on and how to ask interesting and important questions—even during the talks, he asks simple yet deep questions every single time, which I hope I can emulate someday. He is a treasure trove of ideas, and irrespective of the project stage, be it the early brainstorming stage or when I am stuck on something or when we are working on a write-up, after meeting him (mostly virtually, unfortunately), I always end up with much better clarity and having many new ideas to work on.

I owe my gratitude to other members of my thesis committee: Anupam Gupta, Ryan O'Donnell, and Nikhil Bansal. Talking to Anupam, with his constant optimism, is always enlightening. I have benefited from his inspiring teaching style, by attending his lectures and by being a TA. I am especially grateful to him for going out of his way to give feedback on how to improve my presentation skills. I had great fun binge-watching Ryan's youtube lectures on topics ranging from analysis of Boolean functions to infinite expander graphs. Reading the lecture notes of a course taught by Ryan and Venkat at the University of Washington on PCPs during my undergrad was one of the earliest experiences that helped develop my interest in theoretical computer science, and I feel privileged to have both Ryan and Venkat on my thesis committee. I am also indebted to Anupam and Ryan for helping me with my Ph.D. requirements. I thank Nikhil for his encouragement and for his comments on an early draft that were greatly helpful. I am thankful to Pravesh Kothari with whom I enjoyed talking about things, research and otherwise, in the last semester. I am also grateful to Deb Cavlovich for all the help with administrative matters.

Most of the work in this thesis is a joint work and I am thankful to my collaborators: Joshua Brakensiek, Sami Davies, Mordecai Golin, Varun Gupta, Ravishankar Krishnaswamy, Janardhan Kulkarni, Amit Kumar, Euiwoong Lee, Jakub Opršal, Arka Ray, Thomas Rothvoss, Janani Sundaresan, Jakub Tarnawski, Yihao Zhang. While not every collaboration has led to concrete results, I have thoroughly enjoyed every discussion and have learned a lot. I also thank Libor Barto, Amey Bhangale, Boris Bukh, Vincent Cohen-Addad, Florian Frick, Abhishek Shetty, and Xinyu Wu for helpful discussions.

For the past few years, Pittsburgh has become a home away from home. Having lived in zero other cities in the US, I still claim that Pittsburgh is one of the best cities in the US, where you can get everything that a big city offers, without actually feeling like you're in one. But the reason it has become home is the people and I am grateful to the friends here in Pittsburgh: Adithya, without whom I am not sure how I would have survived the pandemic; other roommates (not all at once) Rawal, Eyan, Namit; Prashanth, who helped me during the early days when I moved to the US; Sakshi,

who has been incredibly supportive and encouraging (and also introduced me to *The Marvelous Mrs. Maisel*); Abhiram, Andrii, Ankush, Anup, Biswa, Karthik, Leqi, Nirav, Roger, Shilpa, Sitoshna who have been part of my pandemic social bubble (at various points) and have been there for me through the highs and lows of grad school. I also had fun group lunch/dinners with Ainesh, Alperen, Jackson, Jeff, Joao, Jonathan, Madhusudan, Nic, Omar, Peter, Prashanti, Tim, Vijay; with the Microsoft Research gang at CMU including Anish, Chirags (Gupta and Pabbaraju), Don, Nithin, Phillip, Rahul, Raut, Suhas; Biking and tennis with Akarsh, Jatin, Mihir, Sol; Avalon with David, Guillaume, and Marissa.

Before CMU, I spent a wonderful year at Microsoft Research in Bangalore working with Ravishankar Krishnaswamy. I am grateful to Ravi for his continued encouragement and support during grad school. Even before that, during my undergrad, I am grateful to Ajit Diwan and Nutan Limaye for introducing me to discrete math and theoretical computer science research.

Finally, I am thankful to my family: my uncle and aunt in Virginia, who have always been there for me here in the US, and my parents, my sister Navya, and my cousin Tarak for their unconditional love and encouragement.

Contents

- 1 Introduction** **1**
 - 1.1 Promise Constraint Satisfaction Problems (PCSPs) 2
 - 1.2 Multidimensional Packing and Scheduling 4
 - 1.3 Approximate Hypergraph Vertex Cover and generalized Tuza’s conjecture 5
 - 1.4 Scheduling with non-uniform communication delays 6
 - 1.5 Chapter Credits 7
 - 1.6 Organization 7

- I Promise Constraint Satisfaction Problems** **9**

- 2 Promise Constraint Satisfaction Problems: Introduction** **11**
 - 2.1 PCSPs and Polymorphisms. 11
 - 2.2 Label Cover 13

- 3 Conditional dichotomy of Boolean Ordered PCSPs** **15**
 - 3.1 Introduction 15
 - 3.2 Preliminaries 17
 - 3.3 Algorithm when Shapley values are small 19
 - 3.4 Hardness Assuming Rich 2-to-1 Conjecture 21
 - 3.4.1 Shapley value under random 2-to-1 minor 21
 - 3.4.2 Reduction 25
 - 3.5 Adversarial 2-to-1 minor 27

- 4 d -to-1 Hardness of Coloring 3-colorable graphs with $O(1)$ colors** **31**
 - 4.1 Introduction 31
 - 4.2 Preliminaries 33
 - 4.2.1 d -to-1 Conjecture 33
 - 4.2.2 Low degree influences 34
 - 4.3 d -to-1 hardness for 3-colorable graphs 35
 - 4.3.1 Reducing chromatic number to 3 35
 - 4.3.2 A symmetric Markov chain supported on disjoint tuples 36
 - 4.3.3 Proof of Theorem 33 38
 - 4.4 Reducing multigraph (exact) d -to-1 to $(d + 1)$ -to-1 conjecture 40

5	Rainbow coloring hardness via low sensitivity polymorphisms	43
5.1	Introduction	43
5.1.1	Techniques	44
5.1.2	Prior work on rainbow coloring and related problems	46
5.1.3	Outline	47
5.2	Preliminaries	47
5.2.1	Rainbow Coloring PCSP	47
5.2.2	Complexity measures of functions	48
5.3	Polymorphisms	48
5.3.1	Sensitivity vs certificate complexity	48
5.3.2	Low sensitivity polymorphisms of rainbow coloring	49
5.3.3	High sensitivity polymorphism of RAINBOW(7;6;2)	53
5.4	NP-Hardness	54
5.5	Application: Vector Bin Covering	57
5.5.1	Problem overview	57
5.5.2	Hardness of Vector Bin Covering via Rainbow Coloring	58
5.5.3	Proof of Theorem 57	59
5.6	Adding equality constraints	63
6	Robust Algorithms and SDPs for Promise CSPs	65
6.1	Introduction	65
6.1.1	Robust algorithms	67
6.1.2	Unique Games based hardness	69
6.1.3	Minion characterization of basic SDP	71
6.2	Preliminaries	72
6.3	General Observations	74
6.3.1	Basic SDP setup	74
6.3.2	Generic RHS reduction to “not χ ”	75
6.4	Robust Algorithms	75
6.4.1	CMM is a robust algorithm for MAJ	75
6.4.2	Warm-up for AT: Oblivious LP rounding algorithm for OR	79
6.4.3	Algorithm for AT	80
6.4.4	General case for AT	82
6.5	Unique Games based Hardness	84
6.5.1	Sphere Ramsey Theory	86
6.5.2	Absence of sphere coloring	89
6.6	The SDP minion	98
6.6.1	SDP Minion Definition	99
6.6.2	An alternative Basic SDP	100
6.6.3	From minion homomorphism to SDP rounding algorithm	102
6.6.4	From SDP rounding algorithm to minion homomorphism	102
6.7	Missing Proofs	103

7	Revisiting Alphabet Reduction	109
7.1	Introduction	109
7.2	Preliminaries	111
7.2.1	Rectangular relation and the long code	111
7.2.2	Boolean Fourier analysis	112
7.3	Label Cover to CSP	113
7.3.1	Long code test	114
7.3.2	Constraint test	116
7.3.3	The full test	118
7.4	CSP to Label Cover	119
7.5	Derandomization of the gadget decoding	121

II Structured instances 123

8	Multidimensional Packing and Scheduling Problems	125
8.1	Introduction	125
8.1.1	Our Results	126
8.1.2	Related Work	129
8.1.3	Organization	130
8.2	Preliminaries	130
8.3	Vector Bin Packing	132
8.3.1	Packing Dimension	132
8.3.2	Packing Dimension of Simple Bounded Set Families	133
8.3.3	Hardness of Vector Bin Packing	142
8.4	Vector Scheduling	143
8.4.1	Monochromatic Clique	143
8.4.2	From Monochromatic Clique to Vector Scheduling	147
8.4.3	Hardness of Vector Scheduling via Balanced Hypergraph Coloring	148
8.4.4	Proof of Lemma 154	150
8.5	Hardness of simple k -set cover	153
8.6	SDP Relaxation of Monochromatic-Clique	154
8.6.1	Algorithm when $B > \rho \bar{n}$	154
8.6.2	Integrality gap	155

9	Approximate hypergraph vertex cover and generalized Tuza’s conjecture	157
9.1	Introduction	157
9.1.1	Fractional Tuza’s conjecture and the algorithmic hypergraph Turán problem	158
9.1.2	Vertex cover vs. matching and excluded sub-hypergraphs	159
9.1.3	Vertex cover and set cover on simple hypergraphs	160
9.1.4	Other improved hypergraph vertex cover algorithms	161
9.1.5	Open problems	161
9.2	Preliminaries	163
9.3	LP rounding algorithm for AHTP	164

9.3.1	Color-coding based small vertex cover	164
9.3.2	LP rounding based algorithm for AHTP	165
9.3.3	Analysis of the algorithm and proof of Theorem 161	167
9.3.4	$(t; 2)$ version of AHTP	169
9.4	Forbidden sub-hypergraphs and Tuza's conjecture	170
9.4.1	Explicit construction of tent-free hypergraphs	172
9.5	Vertex cover and set cover on simple hypergraphs	173
9.5.1	Vertex cover on simple t -uniform hypergraphs	173
9.5.2	Set Cover on Simple Set Systems	175
10	Scheduling with non-uniform communication delay	179
10.1	Introduction	179
10.1.1	Our Techniques	180
10.1.2	A Brief History of the Communication Delay Problem	182
10.1.3	Discussion and Open Problems	183
10.1.4	Organization	183
10.2	Unique Machine Precedence Constraints Scheduling problem	184
10.3	Hardness of Scheduling With Non-Uniform Communication Delays	184
10.4	Conditional Hardness of Scheduling With Precedence Constraints on Related Machines	186
10.4.1	Hypothesis of [BN15] implies superconstant hardness of the UMPS prob- lem with unit lengths	191
11	Conclusion	193

List of Figures

- 3.1 An illustration of the two step minor approach: Here $f : \{0,1\}^6 \rightarrow \{0,1\}$ is a Boolean function, $f^0 : \{0,1\}^5 \rightarrow \{0,1\}$ is a minor of f with respect to the function $\pi_1 : [6] \rightarrow [5]$ with $\pi_1(i) = \max(i-1, 1)$, and g is a minor of f^0 with respect to the function $\pi_2 : [5] \rightarrow [3]$ with $\pi_2(i) = \lfloor \frac{i+1}{2} \rfloor$ 21
- 8.1 An illustration of a sunflower-bouquet set family. Here, 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_1 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. 135
- 9.1 The 3-tent 159
- 10.1 Role of the UMPS problem in our hardness reduction. 181
- 10.2 Illustration of the reduction from UMPS to non-uniform communication delays. In the communication delay instance on the right, the dashed arrow precedences have communication delay C_1 while the normal arrow precedences have communication delay 0. 185

Chapter 1

Introduction

Starting with the celebrated PCP theorem [Aro+98], the study of the hardness of approximation has played a key role in theory of computing, including results such as Parallel Repetition [Raz98], and the study of Unique Games Conjecture (UGC) [Kho02b]. This has led to optimal inapproximability results for various computational problems such as 3-SAT [Hås01], and more generally, all Constraint Satisfaction Problems [Rag08] (modulo UGC). Despite this progress, our understanding of approximation algorithms for some problems is lacking, for example, approximate graph coloring, and problems that have been well studied in the algorithms community such as clustering, packing, and scheduling problems.

In this thesis, we make progress on these problems by studying computational problems *under a given promise*. There are two different kinds of promises on the instances: first, we are promised that the instance has a solution that satisfies a stronger property while the goal is to find a solution with a weaker property. For example, given a graph that is promised to be 3-colorable, can we color it with 6 colors in polynomial time? This is an example of Promise Constraint Satisfaction Problems (PCSPs) that generalize the classical Constraint Satisfaction Problems (CSPs) by having weak and strong predicate pairs and the goal is to find a solution satisfying the weak predicates under the promise that there is a solution satisfying the strong predicate. PCSPs are a vast generalization of CSPs, capturing problems such as approximate graph and hypergraph coloring, $(2 + \epsilon)$ -SAT [AGH17].

Formally introduced in a work of Austrin, Guruswami, and Hastad [AGH17], there has been a flurry of works on PCSPs, both on the algorithmic and hardness front. In this thesis, we continue this line of works. Regarding specific PCSPs, for the approximate graph and hypergraph coloring, we prove the hardness of $O(1)$ -coloring a 3-colorable graphs [GS20a] under a weaker conjecture, namely the d -to-1 conjecture of Khot [Kho02b]. Furthermore, we prove improved hardness of rainbow coloring of hypergraphs [GS20b], and also apply these hardness results to prove almost optimal hardness results for Vector Scheduling and Vector Bin Covering [San21]. For the Boolean case, we prove a conditional dichotomy result for monotone Boolean PCSPs [BGS21]. We also study *robust* algorithms for PCSPs where the goal is to output a solution satisfying $1 - \epsilon$ fraction of the constraints on instances promised to be $1 - \epsilon$ satisfiable, with $\epsilon \rightarrow 0$ as $\delta \rightarrow 0$.

A different way of studying computational problems under a promise is when the instances

themselves have a strong structural property. As a concrete example, consider the set cover problem where given a set system, the objective is to find the minimum number of sets whose union is the whole universe of elements. There is an n -factor approximation algorithm for the problem, and this is tight [Fei98]. Suppose that the instance has stronger structural property, namely, that any two sets in the family intersect in at most one element i.e., the underlying set system is *simple* – can we get improved algorithms for the problem, or does the same hardness hold? It turns out that the set cover problem on simple set systems is a useful problem to study, with connections to several other problems. We prove a hardness result on the set cover problem on simple set systems and use it to show optimal (up to constants) hardness of approximation for Vector Bin Packing [San21], whose approximability has been open for more than two decades. The set cover problem on simple set systems is also closely related to the Tuza’s conjecture relating to packing and covering of edges of a graph with triangles, and we use the algorithmic ideas used in the set cover problem to prove the fractional version of the generalized Tuza’s conjecture [GS22]. Finally, we use the idea of studying computational problems under promise to resolve the approximability of scheduling with non-uniform communication delays [Dav+22], where we introduce a problem called Unique Machines Precedence constraints Scheduling (UMPS) which is a very structured instance of unrelated machine scheduling with precedence constraints problem.

1.1 Promise Constraint Satisfaction Problems (PCSPs)

As mentioned earlier, Promise Constraint Satisfaction Problems (PCSPs) are a vast generalization of CSPs, where each predicate now has a weak and strong form. The central question in the study of PCSPs is whether there exists a complexity dichotomy of CSPs extends to PCSPs i.e. if every PCSP is either in P or is NP-complete. We first give a historic overview of the CSP dichotomy theorem. The quest for CSP dichotomy started with a result of Schaefer who proved that every Boolean CSP is either in P or is NP-Hard [Sch78]. Feder and Vardi [FV98] conjectured that the same should hold over arbitrary domains as well. They also showed that the then known algorithmic results all follow by the algebraic closure properties of the CSPs.

This notion was formalized by Jeavons, Cohen, and Gyssens [JCG97; Jea98] and other works [BJK05] that crystallized the (*universal*) *algebraic approach* to CSPs. In the algebraic approach, the higher-order closure properties obeyed by the predicates, namely their *polymorphisms*, are studied. A polymorphism is a function that when applied coordinate-wise to arbitrary satisfying assignments to the predicate, is guaranteed to produce an output that satisfies the predicate. For example, consider an arbitrary instance I of the 2-SAT problem over n variables, and suppose that $\mathbf{x}; \mathbf{y}; \mathbf{z} \in \{0,1\}^n$ are three assignments that satisfy all the constraints in I . Now, if we compute $\mathbf{u} \in \{0,1\}^n$ that is obtained by setting $u_i = \text{MAJ}(x_i; y_i; z_i)$ for all $i \in [n]$, the assignment \mathbf{u} also satisfies all the constraints of I . Thus, the majority function on 3 bits is a polymorphism of the 2-SAT CSP. On the other hand, for the 3-SAT problem, it is not hard to prove that the only polymorphisms are the dictator functions. The algebraic approach has been immensely successful and culminated in the recent resolution of the Feder-Vardi conjecture by Bulatov [Bul17] and

Zhuk [Zhu20]. Further, these proofs yield a precise understanding of the mathematical structure underlying efficient algorithms: if the CSP has a “non-trivial” polymorphisms, the CSP is polytime solvable, and otherwise, it is NP-complete.

The algebraic approach is the key tool towards establishing such a dichotomy result even for PCSPs. The Galois correspondence from the CSP world extends to PCSPs, i.e., the polymorphisms fully capture the computational complexity of the underlying PCSP [Pip02a; BG21a]. This has been extended to show that just the identities satisfied by the polymorphisms suffice to capture the computational complexity of the underlying PCSP [Bar+21]. However, the polymorphisms of PCSPs are much richer, and characterizing which polymorphisms lead to algorithms and which ones lead to hardness has been a challenging problem. Conceptually, the principal difficulty is that the polymorphisms for CSPs are closed under composition (hence referred to as *clones*), whereas for PCSPs, this is no longer the case. As a result, even in the Boolean case, we do not have a dichotomy theorem for PCSPs. On the other hand, this richness of PCSPs has motivated a lot of recent works, both understanding fixed template PCSPs such as variants of graph and hypergraph colorings [Bar+21; KO19], variants of Boolean PCSPs [AGH17; BŽ21], understanding the power of various algorithmic techniques for PCSPs [Bra+20; ČŽ22b].

In this thesis, we continue this line of works. In particular, we prove the following results regarding Promise CSPs.

Conditional Dichotomy of Boolean Monotone PCSPs. Towards establishing a potential Boolean PCSP dichotomy, progress has been made by Fíček, Kozik, Olsák, and Stankiewicz [Fic+19], who obtained a dichotomy result when each predicate is symmetric. In this work, we study Boolean PCSPs that contain the *simplest non-symmetric predicate*, $x \neq y$. We call such Boolean PCSPs *Ordered* as we can also view the implication constraint as an ordering requirement $x \leq y$. We show that Ordered Boolean PCSPs exhibit computational dichotomy, assuming the recently introduced rich 2-to-1 conjecture (which is the perfect completeness analog of the Unique Games Conjecture) of Braverman, Khot, and Minzer [BKM21].

d -to-1 hardness of Coloring 3-Colorable graphs with $O(1)$ Colors. Approximate graph coloring is a canonical PCSP, and it is a major open problem to show NP-Hardness of coloring 3-colorable graphs with $O(1)$ colors. In this work, we prove that the d -to-1 conjecture for any fixed d implies the hardness of coloring a 3-colorable graph with C colors for arbitrarily large integers C . Here, the d -to-1 conjecture of Khot [Kho02b] asserts that it is NP-hard to satisfy an ϵ fraction of constraints of a satisfiable d -to-1 Label Cover instance, for arbitrarily small $\epsilon > 0$. Earlier, the hardness of $O(1)$ -coloring a 4-colorable graphs is known [DMR09] under the 2-to-1 conjecture, which is the strongest in the family of d -to-1 conjectures, and the hardness for 3-colorable graphs is known under a certain “fish-shaped” variant of the 2-to-1 conjecture.

Rainbow Coloring Hardness via low sensitivity polymorphisms. Rainbow coloring of hypergraphs is another PCSP that, for most parameters, is harder than graph coloring, and thus, is easier to show unconditional NP-Hardness. Furthermore, it serves as a good testbed to analyze the polymorphisms of a PCSP that is similar to graph coloring. A k -uniform hypergraph is said to be r -rainbow colorable if there is an r -coloring of its vertices such that every hyperedge intersects all r color classes. Given as input such a hypergraph, finding a r -rainbow coloring of it is NP-hard

for all $k \geq 3$ and $r \geq 2$. Therefore, one settles for finding a rainbow coloring with fewer colors (which is an easier task). When $r = k$ (the maximum possible value), i.e., the hypergraph is k -partite, one can efficiently 2-rainbow color the hypergraph, i.e., 2-color its vertices so that there are no monochromatic edges. In this work, we consider the next smaller value of $r = k - 1$, and prove that in this case, it is NP-hard to rainbow color the hypergraph with $q := d^{\frac{k-2}{2}} e$ colors. In particular, for $k \geq 6$, it is NP-hard to 2-color $(k - 1)$ -rainbow colorable k -uniform hypergraphs.

Vector Bin Covering. We use the analysis of polymorphisms of rainbow coloring PCSP to prove the hardness of Vector Bin Covering. Vector Bin Covering is a multidimensional generalization of Bin Covering. In the d -dimensional Vector Bin Covering instance, 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. This problem is introduced by Alon *et al.* [Alo+98] who gave a $O(\log d)$ factor approximation algorithm. On the hardness front, Ray [Ray21] showed that the 2-dimensional Vector Bin Covering problem is hard to approximate within a factor of $\frac{998}{997}$. We show $\frac{\log d}{\log \log d}$ hardness for the problem, almost matching the $O(\log d)$ factor algorithm [Alo+98].

Robust algorithms for Promise CSPs. In this work, we study robust algorithms for PCSPs i.e., algorithms that output a solution satisfying $1 - f(\epsilon)$ fraction of the constraints when the instance is guaranteed to have a solution satisfying $1 - \epsilon$ fraction of constraints, where $f(\epsilon)$ goes to 0 as ϵ goes to 0. We show that if a Boolean folded PCSP contains Majority or Alternating-Threshold¹ family of polymorphisms, then it admits a robust algorithm. We also show Unique Games based hardness of obtaining robust algorithms for a broad class of PCSPs by showing integrality gaps for basic SDP [Rag08].

Revisiting Alphabet reduction. Along with Gap amplification, Alphabet reduction is a key step in Dinur’s celebrated proof [Din07] of the PCP Theorem. The alphabet reduction used in [Din07] is proved via Assignment Testers. In this chapter, we give a simplified proof of alphabet reduction using a direct test, inspired by reductions between binary PCSPs.

In the rest of the three sections, we give an overview of the multidimensional packing problems, generalized Tuza’s conjecture, and scheduling with non-uniform communication delay.

1.2 Multidimensional Packing and Scheduling

Vector Bin Packing is a multidimensional generalization of Bin Packing where 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. Already when $d = 2$, the problem is APX-hard [Woe97; Ray21]. On the algorithmic front, the PTAS for Bin Packing [VL81] 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^{ϵ} shown by [CK04; Chr+17]. When d is

¹Alternate-Threshold (AT) is a signed variant of the Majority function. For an odd integer L , $AT_L(x_1; x_2; \dots; x_L) = 1$, if $x_1 - x_2 + x_3 - \dots + x_L > 0$, and 0 otherwise. If a PCSP contains AT polymorphisms of all odd arities (for example, (1-in-3-SAT, NAE-3-SAT)), then it can be solved in polynomial time [BG21b].

a fixed constant², 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.

We resolve this gap by proving a $(\log d)$ asymptotic hardness of approximation when d is a large constant, matching the $\ln d + O(1)$ approximation algorithms [CK04; BCS09; BEK16], up to constants. We obtain our hardness result via a reduction from the set cover problem on simple bounded set families – where the underlying set family is simple, each set has a bounded size, and each element appears in a bounded number of sets. We use this hardness to obtain the above theorem using a notion of embedding of set systems that we call *packing dimension*.

Vector Scheduling. Vector Scheduling is another well-studied problem for which the hardness of approximate graph coloring can be used to obtain almost optimal inapproximability results. In the d -dimensional 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_1 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 [Im+19]. We show that these are almost tight by proving that the problem has no $(\log d)^{1-\epsilon}$ factor approximation algorithms assuming NP does not have quasipolynomial time algorithms. We also relate the problem to balanced hypergraph coloring PCSP and use this connection to show (albeit weaker) NP-Hardness result.

1.3 Approximate Hypergraph Vertex Cover and generalized Tuza’s conjecture

A famous conjecture of Tuza [Tuz81; Tuz90] states that the minimum number of edges needed to cover all the triangles in a graph is at most twice the maximum number of edge-disjoint triangles. This conjecture was couched in a broader setting by Aharoni and Zerbib [AZ20] who proposed a hypergraph version of this conjecture and also studied its implied fractional versions. We establish the fractional version of the Aharoni-Zerbib conjecture up to lower-order terms. Specifically, we give a factor $t=2 + O\left(\frac{1}{t \log t}\right)$ approximation based on LP rounding for an algorithmic version of the hypergraph Turán problem (AHTP). The objective in AHTP is to pick the smallest collection of $(t - 1)$ -sized subsets of vertices of an input t -uniform hypergraph such that every hyperedge contains one of these subsets.

²The algorithms are now allowed to run in time $n^{f(d)}$, for some function f .

The algorithmic questions arising in the above study can be phrased as instances of vertex cover on *simple* hypergraphs, whose hyperedges can pairwise share at most one vertex. We prove that the trivial factor t approximation for vertex cover is hard to improve for simple t -uniform hypergraphs. However, for set cover on simple n -vertex hypergraphs, the greedy algorithm achieves a factor $(\ln n) = 2$, better than the optimal $\ln n$ factor for general hypergraphs.

1.4 Scheduling with non-uniform communication delays

We study the problem of scheduling jobs with precedence and non-uniform communication delay constraints on identical machines to minimize the makespan objective function. This classic model was first introduced by Rayward-Smith [Ray87] and Papadimitriou and Yannakakis [PY90]. In this problem, we are given a set J of n jobs, where each job j has a processing length $p_j \geq \mathbb{Z}_+$. The jobs need to be scheduled on m identical machines. The jobs have precedence and communication delay constraints, which are given by a partial order \prec . A constraint $j \prec j^0$ encodes that job j^0 can only start after job j is completed. Moreover, if $j \prec j^0$ and j, j^0 are scheduled on different machines, then j^0 can only start executing at least c_{jj^0} time units after j had finished. On the other hand, if j and j^0 are scheduled on the same machine, then j^0 can start executing immediately after j finishes. The goal is to schedule jobs *non-preemptively* to minimize the makespan objective function, which is defined as the completion time of the last job. In a non-preemptive schedule, each job j needs to be assigned to a single machine i and executed during a contiguous time interval of length p_j . In the classical scheduling notation, the problem is denoted by $P | j \text{ prec}; c_{jk} | j C_{\max}$.

The problem has received renewed interest lately in the applied community due to its relevance in data center scheduling problems and large scale training of ML models. [Cho+11; Guo+12; GCL18; Tar+20]. From a theory perspective, very little was known other than the NP-Hardness of the problem. On the algorithmic front, recently, Maiti et al. [Mai+20] and Davies et al. [Dav+20; Dav+21] designed polylogarithmic approximation algorithms for the special case when all the communication delays are *equal*. On the hardness front, it remained an open problem whether the general problem containing arbitrary communication delays (referred to as non-uniform communication delay scheduling problem) has a constant factor approximation algorithm [SW99b; Ban17].

We answer this question in the negative by showing that for every $\epsilon > 0$, the non-uniform communication delay problem $(P | j \text{ prec}; c_{jk} | j C_{\max})$ does not admit a polynomial-time $(\log n)^{1-\epsilon}$ -approximation algorithm assuming $NP \neq ZTIME_{n^{(\log n)^{O(1)}}$. Our proof of the result follows via a scheduling problem that we call Unique Machine Precedence constraints Scheduling (UMPS), which we believe is a fundamental problem on its own.

1.5 Chapter Credits

Chapter 4, Chapter 5 and Chapter 9 are based on joint works [GS20a], [GS20b] and [GS22] respectively, with Venkatesan Guruswami. Chapter 3 and Chapter 6 are based on [BGS21] and an unpublished work respectively, with Joshua Brakensiek and Venkatesan Guruswami. Chapter 7 is based on [GOS20] a joint work with Venkatesan Guruswami and Jakub Opršal. Chapter 8 is based on [San21]. Chapter 10 is based on [Dav+22] a joint work with Sami Davies, Janardhan Kulkarni, Thomas Rothvoss, Jakub Tarnawski and Yihao Zhang.

1.6 Organization

The thesis is divided into two parts.

In the first part, we study Promise CSPs. We start with some formal definitions in Chapter 2. Then, we prove the conditional dichotomy of Boolean Ordered PCSPs in Chapter 3. We study the approximate graph coloring and rainbow coloring problems in Chapter 4 and Chapter 5 respectively. The robust algorithms for PCSPs are studied in Chapter 6. Finally, we study the alphabet reduction in Chapter 7.

In the second part, we study multidimensional packing problems in Chapter 8, generalized Tuza's conjecture in Chapter 9, and scheduling with non-uniform communication delay problem in Chapter 10. We conclude with some open directions in Chapter 11.

Part I

Promise Constraint Satisfaction Problems

Chapter 2

Promise Constraint Satisfaction Problems: Introduction

In this chapter, we introduce Promise Constraint Satisfaction Problems(PCSPs) formally.

2.1 PCSPs and Polymorphisms.

Constraint satisfaction problems (CSP) have played a very influential role in the theory of computation, providing an excellent testbed for the development of both algorithmic and hardness techniques, which then extend to more general settings. A CSP over domain D is specified by a finite collection of predicates over D . Given an input containing n variables with constraints on the variables using these predicates, the objective is to identify if we can assign values from D to the variables that satisfies all the constraints.

Definition 1. (CSP) Given a k -ary relation $A : D^k \rightarrow \{0,1\}$ over a domain D , the Constraint Satisfaction Problem(CSP) associated with the predicate A takes a set of variables $V = \{v_1; v_2; \dots; v_n\}$ as input which are to be assigned values from D . There are m constraints $(e_1; e_2; \dots; e_m)$ each consisting of $e_i = ((e_i)_1; (e_i)_2; \dots; (e_i)_k) \in D^k$ that indicate that the corresponding assignment should belong to A . The objective is to identify if there is an assignment $V \rightarrow D$ that satisfies all the constraints.

In general, we can have multiple relations $A_1; A_2; \dots; A_l$, and different constraints can use different relations. We denote such a CSP by $CSP(A_1; A_2; \dots; A_l)$.

We give some examples of CSPs.

Example 2. Examples of CSPs include

1. 3-SAT: In this problem, the objective is to assign True or False to variables to satisfy a Boolean formula that is a conjunction of clauses each of which is a disjunction of three literals. In the above notation, we have two predicates, the 3-ary predicate x_y_z together with allowing negation of variables i.e., the predicate $y = \bar{x}$.
2. 2-SAT: This corresponds to the Boolean formula satisfiability problem when each clause contains two literals.

3. *3-coloring.* In this problem, given a graph G , the objective is to assign 3 colors to the vertices of G such that every pair of adjacent vertices are assigned different colors. In the above notation, we have a single predicate $A = f(x; y) : x; y \in [3]; x \neq y$.
4. *1-in-3-SAT.* In this problem, we are given a set of constraints involving three variables, and the objective is to assign True or False to the variables such that in constraint, exactly one variable is assigned True. In the above notation, we have a single predicate $A = f(x; y; z) : x; y; z \in \{0; 1\}; x + y + z = 1$.
5. *NAE-3-SAT.* In this problem, we have a set of constraints involving three variables over the Boolean domain, and the objective is to assign True or False such that in each constraint, both True and False appear. The predicate is $A = f(0; 1)^n \wedge f(0; 0; 0); (1; 1; 1)$.

The key computational challenge in the study of CSPs is to characterize the computational complexity of them i.e., identify which CSPs can be solved in polynomial time, which CSPs are NP-Hard, and whether every CSP is in P or is NP-Hard. In the above set of examples, 2-SAT can be solved in polynomial time while the rest of the CSPs are NP-Hard. The formal study of CSPs was initiated by Schaefer [Sch78] in 1978 when he proved that every Boolean CSP is either in P or is NP-Complete. Feder and Vardi [FV98] conjectured that the same should hold over arbitrary domains as well. After a long line of works, the Feder-Vardi conjecture was resolved in the affirmative by Bulatov [Bul17] and Zhuk [Zhu20] independently.

In this thesis, we study Promise Constraint Satisfaction Problems (PCSPs) that vastly generalize CSPs. In the PCSPs, each predicate has a weak and a strong form—given an instance of PCSP containing n variables with the constraints, the goal is to distinguish between the case that the stronger form can be satisfied vs. even the weaker one cannot be satisfied. We formally define Promise Constraint Satisfaction Problems (PCSPs).

Definition 3. (PCSP) In a Promise Constraint Satisfaction Problem $PCSP(\cdot)$ over a pair of domains $D_1; D_2$, we have a set of pairs of relations $\mathcal{C} = f(A_1; B_1); (A_2; B_2); \dots; (A_l; B_l)$ such that for every $i \in [l]$, A_i is a subset of $D_1^{k_i}$ and B_i is a subset of $D_2^{k_i}$. Furthermore, there is a homomorphism $h : D_1 \rightarrow D_2$ such that for all $i \in [l]$ and $x \in D_1^{k_i}$, $x \in A_i$ implies $h(x) \in B_i$. Given a CSP $(A_1; A_2; \dots; A_l)$ instance, the objective is to distinguish between the two cases:

1. There is an assignment to the variables from D_1 that satisfies every constraint when viewed as CSP $(A_1; A_2; \dots; A_l)$.
2. There is no assignment to the variables from D_2 that satisfies every constraint when viewed as CSP $(B_1; B_2; \dots; B_l)$.

We give some examples of PCSPs.

Example 4. Examples of PCSPs:

1. *(c; s)-approximate Graph Coloring.* A classical example of PCSP is the approximate graph coloring, where given a graph G , the goal is to distinguish between the cases that G can be colored with c colors vs. it cannot be colored with s colors for some $c < s$. In the above language, we have predicates $A; B$ where $A = f(x; y) : x; y \in [c]; x \neq y$, and $B = f(x; y) : x; y \in [s]; x \neq y$.
2. *(1-in-3-SAT, NAE-3-SAT).* Given a 1-in-3-SAT instance that is promised to be satisfiable, the objective is to assign 0; 1 values to the variables such that each constraint is satisfied as

an NAE-3-SAT instance, i.e., both 0 and 1 occur in every constraint.

3. $(2 + \epsilon)$ -SAT: Given a CNF formula where each clause has w literals, with the promise that there is an assignment satisfying $d \frac{w}{2} \epsilon$ literals in each clause, the objective is to find a satisfying assignment to the formula.
4. Approximate rainbow coloring of hypergraphs: A k uniform hypergraph is said to be r -rainbow colorable if there exists a coloring to the vertices such that in each edge, all the colors appear. Given a k -uniform hypergraph that is promised to be $(k - 1)$ -rainbow colorable, the objective is to 2-color¹ the hypergraph.

The study of PCSPs was formally initiated by Austrin, Guruswami, and Håstad [AGH17], and since then, there has been a lot of recent interest in PCSPs, including the development of a systematic theory in [BG21a; Bar+21] and leading to breakthroughs in approximate graph coloring [Bar+21; KO19; WŻ20]. Similar to CSPs, the key question in the study of PCSPs is to understand which PCSPs can be solved in polynomial time, which ones are NP-Hard, and if there is a PCSP dichotomy result. Consider the above example of (1-in-3-SAT, NAE-3-SAT). While the individual CSPs, namely 1-in-3-SAT and NAE-3-SAT are both NP-hard, the above PCSP indeed can be solved in polynomial time [BG21a]. On the other hand, $(c; s)$ -approximate graph coloring is much less understood: it's only in a recent breakthrough result [Bar+21] that $(3; 5)$ -approximate graph coloring is shown to be NP-Hard.

As is the case with CSPs, the key tool underlying these results is the universal algebraic framework where polymorphisms of the PCSP are studied. Polymorphisms capture the closure properties of the satisfying solutions to the PCSP. More formally, we can define polymorphisms of a PCSP as follows.

Definition 5. (Polymorphisms) For $PCSP(\Gamma)$ with $\Gamma = f((A_1; B_1); (A_2; B_2); \dots; (A_l; B_l))g$ where for every $i \in [l]$, $A_i : [q_1]^{k_i} \rightarrow \{0, 1\}g; B_i : [q_2]^{k_i} \rightarrow \{0, 1\}g$, a polymorphism of arity n is a function $f : [q_1]^n \rightarrow [q_2]$ that satisfies the below property for all $i \in [l]$. For all $(v_1; v_2; \dots; v_{k_i})$ such that for all $j \in [k_i]; ((v_1)_j; (v_2)_j; \dots; (v_{k_i})_j) \in A_i$, we have

$$(f(v_1); f(v_2); \dots; f(v_{k_i})) \in B_i$$

We use $\text{Pol}(\Gamma)$ to denote the family of all the polymorphisms of $PCSP(\Gamma)$.

We refer the reader to [Bar+21] for an extensive introduction to PCSPs and polymorphisms.

2.2 Label Cover

We now formally define the Label Cover problem that serves as a starting point for most of our hardness results.

Definition 6. (Label Cover) In the Label Cover instance, we are given a tuple $G = ((V; E); R; \sigma)$ where

1. $(V; E)$ is a graph on vertex set V with edge set E .
2. Each vertex in V has to be assigned a label from the set $\Sigma = [R] = \{r_1; r_2; \dots; r_g\}$.

¹Rainbow 2-coloring is the same as standard 2-coloring of hypergraphs.

3. For every edge $e = (u; v) \in E$, there is an associated relation R_e . This corresponds to a constraint between u and v .

A labeling $\ell : V \rightarrow \Sigma$ satisfies a constraint associated with the edge $e = (u; v)$ if and only if $(\ell(u); \ell(v)) \in R_e$. Given such an instance, the goal is to distinguish if there is a labeling that can satisfy all the constraints or if no labeling can satisfy a significant fraction of constraints.

The PCP theorem [AS98; Aro+98] together with Raz's parallel repetition theorem [Raz98] implies that for every constant $\epsilon > 0$, it is NP-hard to distinguish between the case that a given Label Cover instance has a labeling that satisfies all the constraints vs. no labeling can satisfy more than ϵ fraction of the constraints. This hardness result for Label Cover has been instrumental in showing numerous strong, and sometimes optimal, inapproximability results for various computational problems.

Chapter 3

Conditional dichotomy of Boolean Ordered PCSPs

3.1 Introduction

Towards establishing a potential Boolean PCSP dichotomy, progress has been made by Ficak, Kozik, Olsák and Stankiewicz [Fic+19], who obtained a dichotomy result when each predicate is symmetric. In this chapter, we study Boolean PCSPs that contain the *simplest non-symmetric predicate*, $x \neq y$. We call such Boolean PCSPs *Ordered* as we can also view the implication constraint as an ordering requirement $x \leq y$.

Ordered Boolean PCSPs have come under recent study. The work of Petr [Pet20] (inspired by work of Barto [Bar18b; Bar18a]) considered a special class of Ordered Boolean PCSPs which have an additional predicate $x \neq y$ (this corresponds to allowing negations in the constraints) as well as the requirement that the majority on three bits is *not* a polymorphism. In this setting Petr was able to show that such Ordered Boolean PCSPs are NP-hard. However, the approach considered does not seem immediately extendable to analyzing general Ordered Boolean PCSPs [Bar18a].

The main motivation for studying these PCSPs comes from the fact that adding the additional $x \neq y$ predicate is equivalent to restricting the polymorphisms of the PCSPs to be *monotone functions*. Monotonicity is an influential theme in the study of Boolean functions and complexity theory, and understanding the structure of polymorphisms in the monotone case is an important (and certainly necessary) subcase towards a general characterization of polymorphisms vs. tractability for arbitrary Boolean PCSPs. For the special case of Boolean Ordered PCSPs which include negation constraints, it was conjectured in [Bar18a] that polynomial time tractability is characterized by the existence of majority polymorphisms of arbitrarily large arity.

Our main result is that Boolean Ordered PCSPs exhibit a dichotomy, under the recently introduced *Rich 2-to-1 Conjecture* of Braverman, Khot, and Minzer [BKM21].

Theorem 7. *Assuming the Rich 2-to-1 Conjecture, every Ordered Boolean PCSP is either in P or is NP-Complete. Furthermore, an Ordered PCSP Γ is in P if and only if for every $\epsilon > 0$, there are polymorphisms of Γ with every coordinate having Shapley value at most ϵ . Equivalently, Γ is in P if and only if it has threshold polymorphisms of arbitrarily large arity.*

As a concrete example, recall the PCSP (1-in-3-SAT, NAE-3-SAT) defined in Chapter 2. As it has threshold polymorphisms of arbitrarily large arity, it remains polynomial time solvable even after adding the predicate $x \neq y$. However, if we also add another two-variable predicate $x \notin y$, the PCSP no longer has threshold polymorphisms, and by our above result, it becomes NP-Complete.

We obtain the conditional dichotomy result by analyzing the polymorphisms of the Ordered PCSPs. The key idea in the algebraic approach to PCSPs is that the PCSP is tractable if the polymorphisms are close to symmetric, and the PCSP is hard if all the polymorphisms have a small number of “important” coordinates. More concretely, on the algorithmic front, it has been proved that symmetric polymorphisms of arbitrarily large arities lead to polynomial time algorithms for PCSPs [Bra+20]. On the hardness side, if all the polymorphisms depend on a bounded number of coordinates, then the underlying PCSP is NP-hard [AGH17]. This has been extended to various other notions, including combinatorial ones such as C -fixing [BG16], and topological ones such as having a bounded number of coordinates with non-zero winding number [KO19]. In this work, we study the monotone polymorphisms using analytical techniques.

In particular, we use Shapley value to analyze the monotone polymorphisms. For a monotone function $f : \{0,1\}^n \rightarrow \{0,1\}$, Shapley value of a coordinate i is the probability that on a random path from $(0,0,\dots,0)$ to $(1,1,\dots,1)$, the function value turns from 0 to 1 when we switch the i th coordinate to 1. Initially studied to understand the power of an individual in voting systems [SS54], Shapley value has now found applications in various settings, especially in game theory [Mic+13; NN11]. In our setting, there are two advantages of using Shapley value to study the polymorphisms. First, it is a relative measure of the importance of a coordinate, as opposed to other notions of Influence which are absolute. This helps in bounding the number of coordinates with Shapley value above a certain threshold. Second, it is a versatile measure with combinatorial and analytical interpretations [DDS17] which helps in proving that Shapley value stays consistent under function minors¹, a key property necessary in both the algorithm and the hardness.

Algorithm Overview. We obtain our algorithmic result by using the Basic Linear Programming with Affine relaxation (BLP+Affine relaxation), combined with a structural result regarding the monotone functions with bounded Shapley value. As mentioned earlier, PCSPs with symmetric polymorphisms of arbitrarily large arities can be solved in polynomial time using the BLP+Affine relaxation algorithm [Bra+20]. Our main structural result is that Boolean functions with bounded Shapley value have arbitrarily large threshold functions as minors. Since the set of polymorphisms of a PCSP are closed under taking minors, this proves that the underlying PCSP has arbitrarily large threshold functions as polymorphisms, which then implies that it is in P. The key tool underlying our structural result is a result of Kalai [Kal04] that states that under certain conditions, monotone Boolean functions with arbitrarily small Shapley value have a sharp threshold.

Hardness Overview. We obtain our hardness result assuming the Rich 2-to-1 Conjecture. Braverman, Khot, and Minzer [BKM21] introduced the conjecture as a perfect completeness surrogate of the well known Unique Games Conjecture [Kho02b]. They also proved that the conjecture is

¹A minor (formally defined in Section 3.2) of a function $f : \{0,1\}^n \rightarrow \{0,1\}$ is a function $g : \{0,1\}^m \rightarrow \{0,1\}$ of smaller arity $n > m$ obtained from f by identifying sets of variables together.

equivalent to Unique Games Conjecture when we relax the perfect completeness requirement. The reduction from the Rich 2-to-1 Conjecture to PCSPs follows using the standard Label Cover-Long Code paradigm. The key ingredient in this reduction is a decoding of the Long Codes to a bounded number of coordinates that is consistent under function minors. We decode each Long Code function to the coordinates with (1) Shapley value—as the sum of Shapley values of all the coordinates of any monotone function is equal to 1, there is a bounded number of such coordinates. We argue about the consistency of this decoding using a structural result that states that under a uniformly random minor, Shapley value is roughly preserved.

On the necessity of “richness” in 2-to-1 Conjecture. A natural question is whether our hardness result can be obtained using a weaker assumption such as the 2-to-1 conjecture (whose imperfect completeness version was recently established [KMS17; Din+18a; Din+18b; KMS18]). We shed some light on this question by showing that there are monotone Boolean functions $f : \{0,1\}^{2^n} \rightarrow \{0,1\}$ and $g : \{0,1\}^{2^n} \rightarrow \{0,1\}$ such that g is a minor of f with respect to the 2-to-1 function, both the functions f and g have exactly one coordinate i_1, i_2 respectively, with (1) Shapley value, and yet $i_1 \notin i_2$. Such an adversarial example is interesting from two angles: first, it shows that even using the 2-to-1 conjecture, the Shapley value based decoding is not consistent. Second, it gives an example of agents pairing up maliciously to completely alter the Shapley value. The underlying phenomenon is that the rich 2-to-1 games have “subcode-covering” property, which is absent in the standard 2-to-1 games, helping in preserving the consistency of any biased influence measure such as the Shapley value.

Organization. In Section 3.2, we formally define PCSPs, polymorphisms, and Shapley value. We present the algorithmic and hardness parts of our dichotomy result in Section 3.3 and Section 3.4 respectively. We present the adversarial example of a 2-to-1 minor that alters the Shapley value in Section 3.5.

3.2 Preliminaries

Notations. We use $[n]$ to denote the set $\{1, 2, \dots, n\}$. For a k -ary relation $A \subseteq [q]^k$, we abuse the notation and use A both as a subset of $[q]^k$, and also as a predicate $A : [q]^k \rightarrow \{0,1\}$. For a vector $\mathbf{x} = (x_1, x_2, \dots, x_n) \in \{0,1\}^n$, we use $\text{hw}(\mathbf{x})$ to denote $\sum_{i=1}^n x_i$. For two vectors $\mathbf{x}, \mathbf{y} \in \{0,1\}^n$, we say that $\mathbf{x} \leq \mathbf{y}$ if $x_i \leq y_i$ for all $i \in [n]$. A Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$ is called monotone if $f(\mathbf{x}) \leq f(\mathbf{y})$ for all $\mathbf{x} \leq \mathbf{y}$.

Boolean Ordered PCSPs and Minors of functions. We formally define Boolean Ordered PCSPs.

Definition 8. (Boolean Ordered PCSP) A PCSP $\text{PCSP}(\cdot, \cdot)$ over a pair of domains D_1, D_2 with the set of pairs of relations $\{(A_1, B_1), (A_2, B_2), \dots, (A_l, B_l)\}$ is said to be Boolean Ordered if the following hold.

1. The domains are both Boolean i.e., $D_1 = D_2 = \{0,1\}$.
2. There exists $i \in [l]$ such that $A_i = B_i = \{(0,0), (0,1), (1,1)\}$.

A crucial property satisfied by the polymorphisms of a PCSP Γ , $\text{Pol}(\Gamma)$ is that the family of functions is closed under taking minors. We first define the minor of a function formally.

Definition 9. (*Minor of a function*) For a Boolean function $f : [q]^n \rightarrow [q]$, the function $g : [q]^m \rightarrow [q]$ is said to be a minor of f with respect to the function $\pi : [n] \rightarrow [m]$ if

$$g(x_1; x_2; \dots; x_m) = f(x_{(1)}; x_{(2)}; \dots; x_{(n)}) \quad \forall x_1; x_2; \dots; x_m \in [q]$$

We say that a function g is a minor of f if there exists some π such that g is a minor of f with respect to π .

We are often interested in 2-to-1 minors. A function g is said to be a 2-to-1 minor of f if there exists a 2-to-1 function π such that g is a minor of f with respect to π , where 2-to-1 function is defined below.

Definition 10. (*2-to-1 function*) A function $\pi : [2n] \rightarrow [n]$ is said to be a 2-to-1 function if

$$\pi^{-1}(i) = \{2i-1, 2i\} \subseteq [2n]$$

We use $F_{2 \rightarrow 1}(n)$ to denote the set of all the 2-to-1 functions from $[2n]$ to $[n]$.

By the definition of the polymorphisms, we can infer that if $f \in \text{Pol}(\Gamma)$ for a PCSP Γ , then for all functions g such that g is a minor of f , we have $g \in \text{Pol}(\Gamma)$. Such a family of functions that is closed under taking minors is called as a *minion*. We often refer to the family of polymorphisms of a PCSP as the polymorphism minion.

Shapley value. Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a monotone Boolean function. We can view the monotone Boolean function f as a voting scheme between two parties, and n agents: the winner of the voting scheme when the i th agent votes for $\mathbf{x} \in \{0,1\}^n$ is $f(\mathbf{x})$. The relative power of an agent in a voting scheme is typically measured using the Shapley-Shubix Index, also known as Shapley Value.

Informally speaking, the Shapley Value of a coordinate i is the probability that the i th agent is the altering vote when we start with all zeroes and flip the votes in a uniformly random order. More formally,

Definition 11. (*Shapley value*) Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a monotone Boolean function. Let $\pi \in S_n$ be a uniformly random permutation of $[n]$. For an integer $j \in [n]$, let P_j denote the set of first j elements of π i.e., $P_j := \{\pi(1); \pi(2); \dots; \pi(j)\}$. The Shapley value $\phi_f(i)$ of the coordinate $i \in [n]$ is defined as

$$\phi_f(i) := \Pr_{\pi \in S_n} : (i \in P_{j-1}) \wedge (i \notin P_j) \wedge f(P_{j-1}) = 0 \wedge f(P_j) = 1$$

We also give an alternate definition of Shapley value using the notion of boundary of a coordinate. For a monotone Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$ and coordinate $i \in [n]$, let $B_f(i)$ denote the boundary of the coordinate i i.e.

$$B_f(i) := \{S \subseteq [n] : f(S \cup \{i\}) = 1 \wedge f(S) = 0\}$$

By the monotonicity of f , we can infer that $B_f(i)$ satisfies the following sandwich property that will be useful later.

Proposition 12. Let $f : \{0,1\}^n \rightarrow \{0,1\}$ be a monotone Boolean function and let $i \in [n]$. Then, for every pair of sets $S_1, S_2 \in B_f(i)$ with $S_1 \subseteq S_2$, we have $S \in B_f(i)$ for all S such that $S_1 \subseteq S \subseteq S_2$.

Proof. By the monotonicity of f , we have $f(S \setminus \{i\}) = f(S_1 \setminus \{i\}) = 1$, and thus, $f(S \setminus \{i\}) = 1$. Similarly, we have $f(S \cup \{i\}) = f(S_2 \cup \{i\}) = 0$, and thus, $f(S) = 0$. \square

For an index $j \in \{0,1,\dots,n-1\}$, let $f(j)^{(i)}$ denote the fraction of subsets of $[n]$ of size j that are in $B_f(i)$ i.e.

$$f(j)^{(i)} := \frac{|B_f(i) \cap \binom{[n]}{j}|}{\binom{n}{j}}.$$

We can rewrite the definition of Shapley value of the i th coordinate as the following [Web77]:

$$f(i) = \frac{\sum_{j=0}^{n-1} \binom{n-1}{j} f(j)^{(i)}}{n}. \quad (3.1)$$

3.3 Algorithm when Shapley values are small

In this section, we show that monotone Boolean functions where each coordinate has bounded Shapley value has arbitrarily large threshold functions as minors, thereby proving the algorithmic part of our dichotomy result.

Let L be a positive integer and $0 \leq t \leq L$ be a non-negative integer. We let $\text{THR}_{L,t} : \{0,1\}^L \rightarrow \{0,1\}$ be the threshold function on L variables with threshold t . More formally,

$$\text{THR}_{L,t}(\mathbf{x}) := \begin{cases} 1 & \text{if } \text{hw}(\mathbf{x}) \geq t \\ 0 & \text{otherwise.} \end{cases}$$

For a monotone Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$ and real number $p \in [0,1]$, let $P_p(f)$ denote the expected value of $f(x)$ where each element $x_i; i \in [n]$ is independently set to be 1 with probability p and 0 with probability $1-p$. For every monotone function f , the function $P_p(f)$ is a strictly monotone continuous function in p on the interval $[0,1]$. The value $p_c = p_c(f)$ at which $P_{p_c}(f) = \frac{1}{2}$ is called the *critical probability* of f .

Using the Russo-Margulis Lemma [Rus82; Mar74] and Poincaré Inequality, we can show the following lemma that we need later.

Lemma 13 (Exercise 8.29(e) in [ODO14]). *Let f be a non-constant monotone Boolean function with critical probability $p_c \leq \frac{1}{2}$. Let $p_1 := \frac{1}{(2\epsilon)^2} p_c$ for $\epsilon > 0$. If $p_1 \leq \frac{1}{2}$, then $P_{p_1}(f) \geq \frac{1}{2} - \epsilon$.*

We now define the threshold interval of f .

Definition 14. For a monotone function f and $0 < \epsilon < \frac{1}{2}$, we define $T(f) := [p_2, p_1]$; where p_2 and p_1 are such that $P_{p_1}(f) = \frac{1}{2} + \epsilon$; $P_{p_2}(f) = \frac{1}{2} - \epsilon$.

Kalai [Kal04] proved the following result regarding monotone Boolean functions.

Theorem 15. For every $\epsilon; \delta > 0$, there exists $\epsilon(\epsilon, \delta) > 0$ such that for every monotone Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$ with $f(i) \leq \epsilon$ for all $i \in [n]$ and a $p_c(f) \leq 1 - \delta$, then $T(f) \subseteq [p_c(f) - \epsilon, p_c(f) + \delta]$.

We will use this result to show that for every monotone function where each coordinate has bounded Shapley value has arbitrarily large threshold functions as minor.

Lemma 16. *For every $L \geq 2$, there exists a $\delta := \delta(L) > 0$ such that the following holds. For any monotone Boolean function $f : \{0,1\}^n \rightarrow \{0,1\}$ with*

$$\rho_c(f) \leq \delta \quad \forall i \in [n]$$

there exists a positive integer $L^0 \geq fL; L + 1g$ and a non-negative integer ℓ such that $\text{THR}_{L^0, \ell}$ is a minor of f .

Proof. We obtain $\delta := \delta(L) > 0$ from Theorem 15 by setting $\epsilon = \frac{1}{2^{L+1}}$; $\delta = a = \frac{1}{L^3}$. Our goal is to show that for this parameter δ , for every monotone Boolean function f with each coordinate having Shapley value at most δ , there exists $L^0 \geq fL; L + 1g$ and ℓ such that $\text{THR}_{L^0, \ell}$ is a minor of f .

We assume that f is a non-constant function, else we have a trivial minor by setting $\ell = 0$ or $\ell = L^0$. Let p_c be the critical probability of f .

Case 1: $p_c < a = \frac{1}{L^3}$.

Let $p_1 = L^2 p_c < \frac{1}{L}$. Using Lemma 13, we can conclude that $P_{p_1}(f) \geq 1 - \frac{1}{2L}$. As $P_p(f)$ is monotone, we get that $P_{\frac{1}{L}}(f) > 1 - \frac{1}{2L}$. We let $g : \{0,1\}^{L^0} \rightarrow \{0,1\}$ be a uniformly random minor of f i.e. we choose the function $g : [n] \rightarrow [L]$ by choosing each value $g(i)$ uniformly and independently at random from $[L]$, and we let g to be the minor of f with respect to g .

Note that for every $i \in [L]$, the distribution of $g(fig)$ over the random minor g is the same as sampling a random input to f where we set each bit to 1 with probability $\frac{1}{L}$. As $P_{\frac{1}{L}}(f) \geq 1 - \frac{1}{2L}$, we get that for each $i \in [L]$, $g(fig) = 1$ with probability at least $1 - \frac{1}{2L}$. By union bound, with probability at least $\frac{1}{2}$, $g(fig) = 1$ for all $i \in [L]$. As $f(0; 0; \dots; 0) = 0$, $g(\ell) = 0$ as well. Thus, with probability at least $\frac{1}{2}$, $g = \text{THR}_{L,1}$. Hence, $\text{THR}_{L,1}$ is a minor of f .

Case 2: $p_c > 1 - a = 1 - \frac{1}{L^3}$.

Let f^y be the Boolean dual of f defined as $f^y(x) = 1 - f(x)$. Note that $P_p(f^y) = 1 - P_{1-p}(f)$ for all $p \in [0,1]$. Thus, $p_c(f^y) = 1 - p_c < a$. Using the previous case, we can infer that $\text{THR}_{L,1}$ is a minor of f^y with respect to a function $g : [n] \rightarrow [L]$. The same function g proves that $\text{THR}_{L,1}^y = \text{THR}_{L,L}$ is a minor of f .

Case 3: $a - p_c \leq 1 - a$.

Using Theorem 15, we obtain p_1 such that $P_{p_1}(f) \geq \frac{1}{2}$, and $P_{p_1+\epsilon}(f) \geq 1 - \frac{1}{L^0}$, where $\epsilon = \frac{1}{2^{L+1}}$; $\delta = \frac{1}{L^3}$. As $\delta < \frac{1}{L(L+1)}$, there exists $L^0 \geq fL; L + 1g$ and $\ell \in [L^0]$ such that $p_1 + \epsilon < \frac{1}{L^0}$ and $p_1 > \frac{1}{L^0}$. Thus, we get that $P_{\frac{1}{L^0}}(f) > \frac{1}{2}$ and $P_{\frac{1}{L^0}+\epsilon}(f) < \frac{1}{2}$. Let $g : \{0,1\}^{L^0} \rightarrow \{0,1\}$ be a uniformly random minor of f i.e. we choose $g : [n] \rightarrow [L^0]$ by setting each value uniformly and independently at random from $[L^0]$ and set g to be the minor of f with respect to g . For a vector $\mathbf{x} \in \{0,1\}^{L^0}$ with $\text{hw}(\mathbf{x}) = \frac{1}{2}$, with probability greater than $1 - \frac{1}{2^{L+1}}$, $g(\mathbf{x}) = 1$. Similarly, for $\mathbf{x} \in \{0,1\}^{L^0}$ with $\text{hw}(\mathbf{x}) = \frac{1}{2} - \epsilon$, with probability greater than $1 - \frac{1}{2^{L+1}}$, $g(\mathbf{x}) = 0$. Thus, with non-zero probability, $g(\mathbf{x}) = 1$ for all $\mathbf{x} \in \{0,1\}^{L^0}$ with $\text{hw}(\mathbf{x}) = \frac{1}{2}$ and $g(\mathbf{x}) = 0$ for all $\mathbf{x} \in \{0,1\}^{L^0}$ with $\text{hw}(\mathbf{x}) = \frac{1}{2} - \epsilon$. In other words, with non-zero probability, g is equal to $\text{THR}_{L^0, \ell}$. Thus, $\text{THR}_{L^0, \ell}$ is a minor of f . \square




Figure 3.1: An illustration of the two step minor approach: Here $f : \{0,1\}^6 \rightarrow \{0,1\}$ is a Boolean function, $f^d : \{0,1\}^5 \rightarrow \{0,1\}$ is a minor of f with respect to the function $\pi_1 : [6] \rightarrow [5]$ with $\pi_1(i) = \max(i-1, 1)$, and g is a minor of f^d with respect to the function $\pi_2 : [5] \rightarrow [3]$ with $\pi_2(i) = \lfloor \frac{i+1}{2} \rfloor$.

Using the existence of arbitrarily large arity threshold minors, the algorithmic part of our Dichotomy result follows immediately.

Theorem 17. *Let Γ be a Promise CSP template. Suppose that for every $\epsilon > 0$, there exists a function $f \in \text{Pol}(\Gamma)$; $f : \{0,1\}^n \rightarrow \{0,1\}$ such that $\text{thr}_i(f) \geq \epsilon$ for all $i \in [n]$. Then, $\text{PCSP}(\Gamma) \in \text{P}$.*

Proof. Using Lemma 16, we can conclude that there are infinitely many positive integers L such that there exists $\pi \in \{0,1\}^{L \times L}$ with $\text{THR}_L(\pi) \in \text{Pol}(\Gamma)$. As the threshold functions are symmetric², $\text{Pol}(\Gamma)$ has symmetric polymorphisms of infinitely many arities. Thus, using the BLP+Affine algorithm of [Bra+20], $\text{PCSP}(\Gamma)$ can be solved in polynomial time. \square

We remark that the above result is inspired by a special case shown by Barto [Bar18b] that a Boolean Ordered PCSP is polytime tractable if it has cyclic polymorphisms of arbitrarily large arities.

3.4 Hardness Assuming Rich 2-to-1 Conjecture

In this section, we prove the hardness part of our dichotomy result. First, we prove that Shapley value is preserved under uniformly random 2-to-1 minors, and then we use this to show the hardness assuming the Rich 2-to-1 Conjecture.

3.4.1 Shapley value under random 2-to-1 minor

Let $f : \{0,1\}^{2n} \rightarrow \{0,1\}$ be a monotone Boolean function with $\text{thr}_1(f) \geq \epsilon$ for some absolute constant $\epsilon > 0$. Let $g : \{0,1\}^n \rightarrow \{0,1\}$ be a minor of f with respect to the uniformly random 2-to-1 function $\pi : [2n] \rightarrow [n]$. Our goal in this subsection is to show that $\mathbb{E}[\text{thr}_1(g(\pi))]$ is bounded away from 0 for some function $\epsilon := \epsilon(\epsilon) > 0$. We prove this in two steps. (See Figure 3.1)

²A function $f : \{0,1\}^n \rightarrow \{0,1\}$ is said to be symmetric if it is unchanged by any permutation of the input variables.

1. First, we consider the minor of f , $f^0 : \{0,1\}^{2^n-1} \rightarrow \{0,1\}$ obtained with respect to $\pi_1 : [2n] \rightarrow [2n-1]$ where $\pi_1(1) = \pi_1(2) = 1$; $\pi_1(i) = i-1$ if $i \geq 3$; $\pi_1(4) = \dots = \pi_1(2n) = 2n$. We show that $f^0(1) = \frac{1}{2}$.
2. Next, we consider a minor g of f^0 obtained with respect to the function $\pi_2 : [2n-1] \rightarrow [n]$ which has $\pi_2(1) = 1$ while the rest $2n-2$ values are chosen using a uniformly random partition of $[2n-2]$ into $n-1$ pairs. We show that $E_{\pi_2}[g(1)] = \frac{1}{2}$ for some function $g := (\cdot)$.

Note that the process of first taking the f^0 minor and then obtaining g by partitioning $[2n-2]$ into $n-1$ uniformly random pairs is equivalent to taking a uniformly random 2-to-1 minor of f . Thus, the two steps together prove the required Shapley value property of the uniformly random 2-to-1 minor.

The first step is captured by the following lemma.

Lemma 18. *Let $f : \{0,1\}^{2^n} \rightarrow \{0,1\}$ and $f^0 : \{0,1\}^{2^n-1} \rightarrow \{0,1\}$ be monotone Boolean functions such that f^0 is a minor of f with respect to the function $\pi_1 : [2n] \rightarrow [2n-1]$ defined as $\pi_1(i) = \max(i-1, 1)$. If $f(1) = \frac{1}{2}$, then $f^0(1) = \frac{1}{2}$.*

Proof. We recall a bit of notation: let $B_f(1)$ denote the boundary of the coordinate 1 in the function f i.e. the family of all the sets $S \subseteq [2n]$ such that $f(S) = 0$; $f(S \cup \{1\}) = 1$. For an integer $j \in \{0,1,\dots,2n-1\}$, let $\mu_f(j)$ denote the fraction of subsets of $[2n]$ of size j that are in $B_f(1)$. For ease of notation, we let $\mu(j) = \mu_f(j)$, and $\mu^0(j) = \mu_{f^0}(j)$. Consider a set $S \subseteq [2n]$ such that $S \in B_f(1)$. Note that

$$S^0 = \{i \in [2n-1] : i \geq 2; i \in S\}$$

satisfies $S^0 \in B_{f^0}(1)$. Suppose that $S_1, S_2 \in B_f(1)$ such that $|S_1 \cap S_2| = j$, $S_1 \not\subseteq S_2$ and $2 \leq |S_1 \setminus S_2|$. Then, the above definition satisfies $S_1^0 \not\subseteq S_2^0$, $S_1^0 \cap S_2^0 \in B_{f^0}(1)$ and $|S_1^0 \cap S_2^0| = j$. This implies that

$$\mu_f(S : S \in B_f(1); |S \cap S^c| = j; 2 \leq |S \setminus S^c|) = \mu_{f^0}(1) \setminus \binom{[2n-1]}{j}$$

Similarly,

$$\mu_f(S : S \in B_f(1); |S \cap S^c| = j; 2 \leq |S \setminus S^c|) = \mu_{f^0}(1) \setminus \binom{[2n-1]}{j-1}$$

Summing the two, we obtain that

$$\mu_f(S : S \in B_f(1); |S \cap S^c| = j) = \mu_{f^0}(1) \setminus \binom{[2n-1]}{j-1} + \mu_{f^0}(1) \setminus \binom{[2n-1]}{j}$$

We can rewrite it as

$$\binom{2n-1}{j} \mu(j) = \binom{2n-2}{j} \mu^0(j) + \binom{2n-2}{j-1} \mu^0(j-1) \quad \forall j \in [2n-2]$$

As $\binom{2n-1}{j} = \binom{2n-2}{j} + \binom{2n-2}{j-1}$ for every $j \in [2n-2]$, we get that

$$\binom{2n-1}{j} = \binom{2n-2}{j} + \binom{2n-2}{j-1}$$

for all $j \in [2n-2]$. Also note that $\binom{2n-1}{0} = \binom{2n-2}{0}$, and $\binom{2n-1}{2n-1} = \binom{2n-2}{2n-2}$. Summing over all these inequalities, we get that

$$\sum_{j \in \{0, 1, \dots, 2n-2\}} \binom{2n-1}{j} = \frac{1}{2} \sum_{j \in \{0, 1, \dots, 2n-1\}} \binom{2n-1}{j} = n$$

Thus,

$$f^0(1) = \frac{\sum_{j \in \{0, 1, \dots, 2n-2\}} \binom{2n-1}{j}}{2n-1} = \frac{1}{2} \quad \square$$

Before proving the second step, we prove the following key lemma regarding the distribution of the boundary subsets.

Lemma 19. Let $f^0 : \{0, 1\}^{[2n-1]} \rightarrow \{0, 1\}$ be a monotone Boolean function such that $f^0(1) = \frac{1}{2}$. For an integer $j \in \{0, 1, \dots, 2n-2\}$, let $\theta(j) = f^0(j)^{(1)}$. Then, there exists an absolute constant $\epsilon := \epsilon(\epsilon) > 0$ such that

$$\frac{\sum_{j=0}^{n-1} \theta(2j)}{n} \geq \frac{1}{2} + \epsilon$$

Proof. We prove that there exist constants (depending on ϵ) $c_1 < c_2, c > 0$ such that for all j such that $c_1 n \leq j \leq c_2 n$, we have $\theta(j) \geq c$, and $c_2 \leq c_1 + \frac{\epsilon}{2}$. This directly implies the required claim with $\epsilon = (c^2)$.

For a pair of integers $0 \leq i < j \leq 2n-2$, we define the following parameter $\theta(i; j)$ as the fraction of the pair of subsets $(S; T)$ where $S; T \in \{0, 1\}^{[2n-1]}$; $jSj = i; jTj = j; S \leq T$ that satisfy $S \leq B_{f^0}(1); T \leq B_{f^0}(1)$.

$$\theta(i; j) = \frac{|\{(S; T) : jSj = i; jTj = j; S \leq T; S \leq B_{f^0}(1); T \leq B_{f^0}(1)\}|}{\binom{2n-2}{i} \binom{2n-2}{j}}$$

We first claim that there exist constants (depending on ϵ) $c_1 < c_2, c > 0$ such that $\theta(c_1 n; c_2 n) \geq c$, and $c_2 \leq c_1 + \frac{\epsilon}{2}$. Consider a uniformly random permutation of $[2n-1]$ denoted by $\pi = (\pi(1); \pi(2); \dots; \pi(2n-2))$. For an integer $j \in \{0, 1, \dots, 2n-2\}$; let S_j be the random variable that is the union of the prefix of π containing the first j elements.

$$S_j := \{\pi(1); \pi(2); \dots; \pi(j)\}; j \in \{0, 1, \dots, 2n-2\}$$

For each $j \in \{0, 1, \dots, 2n-2\}$, the subset S_j is uniformly distributed in $\binom{[2n-1]}{j}$. For $j \in \{0, 1, \dots, 2n-2\}$; let X_j be the indicator random variable for the event that $S_j \leq B_{f^0}(1)$. By the definition of $\theta(j)$, we get

$$E[X_j] = \theta(j) \quad j \in \{0, 1, \dots, 2n-2\}$$

Let $X = X_0 + X_1 + \dots + X_{2n-2}$ be the number of subsets in the set family $(S_0, S_1, S_2, \dots, S_{2n-2} = [2n-1] \setminus \{j\})$ that are in $B_{f^0}(1)$. Using Equation (3.1), we get

$$E[X] = \binom{2n-1}{2}:$$

Using Jensen's inequality, we get that

$$E \left[\frac{X}{2} \right]^2 \leq \frac{\binom{2n-1}{2}}{2} = \frac{1}{2} \binom{2n-1}{2} = \frac{1}{2} (2n-2) + n-1 = \frac{2n-1}{2}$$

wherein the final inequality, we used the fact that $n \geq 1$. Note that for every $i < j$, the marginal distribution of (S_i, S_j) is the uniform distribution over all the pairs of subsets (S, T) where $S, T \subseteq [2; 3; \dots; 2n-1] \setminus \{j\}; |S| = i; |T| = j; S \cap T = \emptyset$. Thus, by the definition of $\theta(i; j)$, we get that $\theta(i; j) = E[X_i X_j]$, for $0 \leq i < j \leq 2n-2$. Therefore we have

$$E \left[\frac{X}{2} \right]^2 = E \left[\prod_{0 \leq i < j \leq 2n-2} X_i X_j \right] = \prod_{0 \leq i < j \leq 2n-2} E[X_i X_j] = \prod_{0 \leq i < j \leq 2n-2} \theta(i; j)$$

Thus,

$$\prod_{0 \leq i < j \leq 2n-2} \theta(i; j) \leq \frac{2n-1}{2}$$

This implies that the expected value (over $i; j$) of $\theta(i; j)$ is at least $\frac{2}{2n-1}$. Thus, with probability (over $i; j$) at least $\frac{2}{4}$, we have $\theta(i; j) \geq \frac{2}{4}$. Hence, there exist integers $p; q$ such that $q \leq p \leq \frac{2}{8}n$ and $\theta(p; q) \geq \frac{2}{4}$, which proves the required claim with $p = c_1 n; q = c_2 n; c = \frac{2}{4}$.

Next, we claim that $\theta(j) \geq c$ for all j such that $c_1 n \leq j \leq c_2 n$. Fix an integer j with $c_1 n \leq j \leq c_2 n$. Consider a uniformly random sequence of subsets $S_1, S_2, S_3 \subseteq [2n-1] \setminus \{j\}$ such that $|S_1| = c_1 n; |S_2| = j; |S_3| = c_2 n$. The probability that $S_1 \subseteq B_{f^0}(1); S_3 \subseteq B_{f^0}(1)$ is equal to $\theta(c_1 n; c_2 n)$ which is at least $\frac{2}{4}$. Thus, using Proposition 12, with probability at least $\frac{2}{4}$, $S_2 \subseteq B_{f^0}(1)$. Note that the distribution of S_2 is uniform in $\binom{[2n-1] \setminus \{j\}}{j}$, and thus, we have $\theta(j) \geq \frac{2}{4}$. \square

We now prove the second step in the proof.

Lemma 20. Suppose that $f^0 : \{0, 1\}^{2n-1}$ is a monotone Boolean function such that $f^0(1) = 1$ with probability $\frac{1}{n}$. Let g be a random minor of f^0 with respect to $\pi : [2n-1] \rightarrow [n]$ which is obtained by setting $\pi(1) = 1$, and for every $i > 1$, we randomly choose $j_1; j_2 \subseteq [2n-1] \setminus \{i\}$ (without replacements) and set $\pi(j_1) = \pi(j_2) = i$. In other words, we choose a uniformly random partition of $[2n-1] \setminus \{1\}$ into $n-1$ pairs $P_2; P_3; \dots; P_n$ and set $\pi(j) = i$ if $j \in P_i$. Then, there exists $\epsilon := \epsilon(n) > 0$ such that

$$E \left[\sum_{g(1)} \right] \geq \epsilon:$$

Proof. For ease of notation, we let $\theta(j) = f^0(j)^{(1)}$ and $g(j) = g(j)^{(1)}$. For a set $S \subseteq [n] \setminus \{1\}$ and a function $\pi : [2n-1] \rightarrow [n]$ with $\pi(1) = 1$, and $\pi^{-1}(i) = \{j_1, j_2\}$ for all $i \in [2; 3; \dots; n]$, let $\pi^{-1}(S)$ be the $2|S|$ sized subset of $[2; 3; \dots; 2n-1]$ defined as follows:

$$\pi^{-1}(S) := \{j \in [2; 3; \dots; 2n-1] : \pi(j) \in S\}$$

For every set $S \subseteq [2n-1] \setminus [n]$ when $\pi : [2n-1] \rightarrow [n]$ is a uniformly random 2-to-1 minor with $\pi(1) = 1$, and the rest $2n-2$ elements are partitioned into $n-1$ pairs uniformly at random, the set $\pi^{-1}(S)$ is distributed uniformly in $\binom{[2n-1] \setminus [n]}{2^j S_j}$. Also note that $S \subseteq B_g(1)$ if and only if $\pi^{-1}(S) \subseteq B_{f \circ \pi}(1)$. Thus, for every set $S \subseteq [2n-1] \setminus [n]$, the probability that $S \subseteq B_g(1)$ (over the choice of π) is equal to $\mathbb{P}(2^j S_j)$. Summing over all such sets of size j , we get that for every $j \subseteq [0; 1; \dots; n-1]_g$, the expected value of $\mathbb{E}_\pi[g(j)]$ is equal to $\mathbb{E}(2^j)$.

$$\mathbb{E}_\pi[g(j)] = \mathbb{E}(2^j) \mathbb{1}_{j \subseteq [0; 1; \dots; n-1]_g}$$

By using Lemma 19, we can infer that there exists $\epsilon = \epsilon(n) > 0$ such that $\mathbb{P}_{j=0}^{n-1} \mathbb{E}_\pi[g(j)] = \sum_{j=0}^{n-1} \mathbb{E}(2^j) \mathbb{1}_{j \subseteq [0; 1; \dots; n-1]_g}$. Using Equation (3.1), we get

$$\mathbb{E}_\pi[g(1)] = \mathbb{E}_\pi \left[\frac{\sum_{j=0}^{n-1} g(j)}{n} \right] = \frac{\sum_{j=0}^{n-1} \mathbb{E}_\pi[g(j)]}{n} : \square$$

Lemma 18 and Lemma 20 together prove that Shapley value behaves well under uniformly random 2-to-1 minors for monotone Boolean functions.

Lemma 21. *Suppose that $f : \{0,1\}^{2n} \rightarrow \{0,1\}$ is a monotone Boolean function such that $f(1) \geq \epsilon$ for some absolute constant $\epsilon > 0$ with $\frac{1}{n}$. Then, there exists $\delta = \delta(\epsilon) > 0$ such that*

$$\mathbb{E}[g(\pi(1))] \geq \delta$$

where g is a minor of f with respect to the uniformly random 2-to-1 function π .

Proof. Combining Lemma 18 and Lemma 20, we can conclude that for every $i \subseteq [2n]; i > 1$, when $\pi : [2n] \rightarrow [n]$ is a uniformly random 2-to-1 minor conditioned on the fact that $\pi(1) = i$, we have $\mathbb{E}[g(\pi(1))] \geq \delta$. Taking average over all the $i \subseteq [2n]; i > 1$, we get a proof that the same inequality holds when π is a uniformly random 2-to-1 minor. \square

3.4.2 Reduction

In his celebrated work proposing the Unique Games Conjecture [Kho02a], Khot also proposed the “2-to-1 conjecture” that the strong hardness of Label Cover holds when all the constraints of the Label Cover are 2-to-1 functions. The imperfect completeness version of this conjecture was recently established in a striking sequence of works [KMS17; Din+18a; Din+18b; KMS18]. Braverman, Khot, and Minzer [BKM21] put forth a stronger conjecture that states that the hardness of Label Cover holds when the distribution of 2-to-1 functions on edges incident on every vertex $u \subseteq L$ is uniform over F_{2^l-1} .

Definition 22. (Rich 2-to-1 Label Cover instances) *We call a Label Cover instance $G = (G; L; R; \pi)$ with $G = (L \cup R; E)$ being a bipartite graph³ a rich 2-to-1 instance if the following hold.*

³In the definition of the Label Cover problem in Chapter 2, we have assigned the same alphabet to all the vertices in a Label Cover instance. Here, in the context of bipartite Label Cover instances, we allow different alphabets to the left and right sides of the graph.

1. There exists an integer k such that $L = [2^k]$, $R = [k]$, and every projection constraint $e, e \in E$ is a 2-to-1 function.
2. For every vertex $u \in L$, the distribution of 2-to-1 functions P_u obtained by first sampling a uniformly random neighbor v of u , and then picking $e: e = (u; v)$, is uniform over $F_{2^k \times 1}(\mathbb{F}_2)$.

Conjecture 23. (Rich 2-to-1 Conjecture with Perfect Completeness) [BKM21] For every $\epsilon > 0$, there exists an integer $k = k(\epsilon)$ such that given a rich 2-to-1 Label Cover instance G , it is NP-Hard to distinguish between the following.

1. There is a labeling that satisfies all the constraints of G .
2. No labeling can satisfy more than ϵ fraction of the constraints of G .

We are now ready to state the hardness part of our dichotomy. It is proved using the Label Cover-Long Code framework. This reduction is standard in the PCSP literature, see e.g., [Bar+21].

Theorem 24. Assume the Rich 2-to-1 Conjecture. Let $PCSP(\mathcal{F})$ be a Boolean Ordered PCSP such that there exists an absolute constant $\epsilon > 0$ with $\max_{i \in [n]} f(i) \geq \epsilon$ for all functions $f: \mathbb{F}_2 \rightarrow \mathbb{F}_2$. Then $PCSP(\mathcal{F})$ is NP-Hard.

Proof. Let $\mathcal{F} = (f(A_1; B_1); (A_2; B_2); \dots; (A_l; B_l))$ be the PCSP under consideration, where each A_i is a subset of $\mathbb{F}_2^{k_i}$ for all $i \in [l]$, and similarly, each B_i is a subset of $\mathbb{F}_2^{k_i}$ for all $i \in [l]$. We start from a rich 2-to-1 Label Cover instance $G = (G; [2^k]; [k]; \mathcal{F})$ with $G = (L \cup R; E)$. For ease of notation, we use L_w to denote L if $w \in L$, and R if $w \in R$. For every vertex $w \in L \cup R$, we have a set of 2^k nodes denoted by $L_w = \{fwg : f \in \mathbb{F}_2^k\}$ referred to as the long code corresponding to w . The elements of our output PCSP instance V is the union of all the long code nodes.

$$V = \bigcup_{w \in L \cup R} L_w$$

We add two types of constraints.

1. Polymorphism Constraints. For every $i \in [l]$, we add the following constraints using the pair of predicates $(A_i; B_i)$. For every $w \in L \cup R$, and multiset of vectors $\mathbf{x}^1; \mathbf{x}^2; \dots; \mathbf{x}^{k_i} \in \mathbb{F}_2^{k_i}$ satisfying

$$(\mathbf{x}^1; \mathbf{x}^2; \dots; \mathbf{x}^{k_i}) \in A_i \text{ } \forall j \in [k_i];$$

we add the constraint on the k_i nodes $\{fw; \mathbf{x}^1g; fw; \mathbf{x}^2g; \dots; fw; \mathbf{x}^{k_i}g\}$.

2. Equality Constraints. For every edge $e = (u; v)$ of the Label Cover instance with the constraint $e: [2^k] \rightarrow [k]$, we add the following set of equality constraints. For every $\mathbf{x} \in \mathbb{F}_2^{k_u}$ and $\mathbf{y} \in \mathbb{F}_2^{k_v}$ such that for all $j \in [2]$, $\mathbf{x}_j = \mathbf{y}_{e(j)}$, we add an equality constraint between $fu; \mathbf{x}g$ and $fv; \mathbf{y}g$ ensuring that the two nodes are assigned the same value. The fact that we can add the equality constraints follows either by identifying the variables together, or by observing that the polymorphism minion of any PCSP remains the same when we add the equality predicate (see Chapter 5).

Completeness. Suppose that there exists a labeling ℓ that satisfies all the constraints of the Label Cover instance. For every node $fw; \mathbf{x}g \in V$, we assign the dictator function $\mathbf{x}_{(w)} \in \mathbb{F}_2^{k_w}$. By

the way we have added the polymorphism constraints, any dictator assignment satisfies them. The equality constraints are also satisfied as the labeling satisfies all the constraints of G .

Soundness. Suppose that there exists an assignment $f : V \rightarrow \{0,1\}^g$ that satisfies all the polymorphism constraints and the equality constraints. Then, we claim that there exists a labeling that satisfies $\epsilon = \epsilon(n) > 0$ fraction of the constraints of the Label Cover instance G .

For a vertex $w \in L \cup R$, let $f_w : \{0,1\}^g \rightarrow \{0,1\}^g$ denote the function f restricted to L_w . Note that f_w is a polymorphism of the PCSP for all $w \in L \cup R$. As every polymorphism of G has a coordinate with Shapley value at least ϵ , for every $u \in L$, we define the set $S(u)$ that is non-empty as follows:

$$S(u) = \{i \in [2^g] : f_u(i) = g\}$$

As $\sum_{i \in [2^g]} f_u(i) = 1$ for all functions $f : \{0,1\}^g \rightarrow \{0,1\}^g$, we have $|S(u)| \geq \epsilon$ for all $u \in L$.

As a corollary of Lemma 21, we can conclude that there exists $\epsilon = \epsilon(n) > 0$ such that for every monotone Boolean function $f : \{0,1\}^{2^g} \rightarrow \{0,1\}^g$ with $f(i) = g$, when g is a minor of f with respect to a uniformly random 2-to-1 function $\pi : [2^g] \rightarrow [g]$, $g(\pi(i)) = \frac{1}{2}$ with probability at least $\frac{\epsilon}{2}$. Note that applying Lemma 21 requires that $\epsilon \geq \frac{1}{2}$. However, even when $\epsilon < \frac{1}{2}$, by picking the coordinate with the largest Shapley value, we can still assume that in every long code function, there is a coordinate with Shapley value at least $\frac{\epsilon}{2} = \epsilon'$, and then apply Lemma 21. Using this ϵ' , for every $v \in R$, we define the set $S(v)$ as

$$S(v) = \{i \in [2^g] : f_v(i) = \frac{g}{2}\}$$

By definition, we have $|S(v)| \geq \epsilon'$ for all $v \in R$. As the Label Cover instance is rich 2-to-1, for every $u \in L$, when we pick a uniformly random edge $e = (u, v)$ adjacent to u with constraint $\pi_e : [2^g] \rightarrow [g]$, with probability at least $\frac{\epsilon}{2}$, there exist $i_1 \in [2^g]; i_2 \in [g]$ such that $f_u(i_1) = g$, $f_v(i_2) = \frac{g}{2}$, and $\pi_e(i_1) = i_2$.

We now pick a labeling σ of G by picking uniformly random label from $S(w)$ for all $w \in L \cup R$. By the above argument, for every $u \in L$, the expected number of constraints of G that are adjacent to u that the labeling σ satisfies is at least $\frac{\epsilon}{2} \cdot \frac{\epsilon}{2}$. Summing over all $u \in L$, σ satisfies at least $\frac{\epsilon^2}{4}$ fraction of the constraints of G in expectation. Thus, there exists a labeling to G that satisfies $\epsilon = \frac{\epsilon^2}{4} > 0$ fraction of the constraints, which completes the proof. \square

3.5 Adversarial 2-to-1 minor

We construct an example of a 2-to-1 minor where the Shapley value alters completely after taking the minor.

Theorem 25. *Let $n \geq 2$ be a positive integer. There exist two monotone Boolean functions $f : \{0,1\}^{2^{2n}} \rightarrow \{0,1\}^g$ and $g : \{0,1\}^{2^n} \rightarrow \{0,1\}^g$ such that g is a 2-to-1 minor of f with respect to the 2-to-1 function $\pi : [2^{2n}] \rightarrow [2^n]$ defined as $\pi(i) = \lfloor \frac{i}{2} \rfloor$. Furthermore,*

1. $f(1) = g(1)$, and $f(j) = o(1)$ for all $j > 1$.
2. $f(3) = g(1)$, and $f(i) = o(1)$ for all $i \in [2^{2n}]; i \notin \{1, 3\}$.

Proof. Similar to the proof of Theorem 24, we construct the minor function pair in two steps.

1. First, we construct Boolean monotone functions $f : \{0,1\}^{2n-1} \rightarrow \{0,1\}$ and $g : \{0,1\}^n \rightarrow \{0,1\}$ such that g is a minor of f with respect to the function $\phi : [2n-1] \rightarrow [n]$ defined as $\phi(1) = 1$; $\phi(i) = \lfloor \frac{i+1}{2} \rfloor$ for all $i > 1$. Furthermore, $g(1) = \phi(1)$, and $g(j) = o(1)$ for all $j > 1$. We also have $f(2) = \phi(1)$, and $f(i) = o(1)$ for all $i \in [2n-1]; i \neq 2$.
2. We define the function $f^\theta : \{0,1\}^{2n} \rightarrow \{0,1\}$ as

$$f^\theta(y_1; y_2; \dots; y_{2n}) = f(y_1; y_3; \dots; y_{2n})$$

Note that g is a minor of f^θ with respect to the 2-to-1 function $\psi : [2n] \rightarrow [n]$ defined as $\psi(i) = \lfloor \frac{i}{2} \rfloor$. Furthermore, by definition, we have $f^\theta(3) = \phi(1)$, and $f^\theta(i) = o(1)$ for all $i \in [2n]; i \neq 3$.

Henceforth, our goal is to construct a pair of functions as in the first step above.

We define a partial Boolean function to be a function from $\{0,1\}^n \rightarrow \{0,1\}^?$. A partial Boolean function on n variables is monotone if for all $\mathbf{p} \geq \mathbf{q}$ and $\mathbf{q} \in \text{dom}(f)$ such that $f(\mathbf{p}) = 1$, then $f(\mathbf{q}) = 1$, and if $f(\mathbf{q}) = 0$, then $f(\mathbf{p}) = 0$.

Now, consider $g : \{0,1\}^n \rightarrow \{0,1\}$ to be

$$g(\mathbf{x}) = \begin{cases} 1 & \text{if } \sum_{j=2}^n x_j \geq \frac{51n}{100} \\ 0 & \text{if } \sum_{j=2}^n x_j \leq \frac{49n}{100} \\ x_1 & \text{if } \frac{49n}{100} < \sum_{j=2}^n x_j < \frac{51n}{100} \end{cases}$$

By definition, g is a monotone function, and using Equation (3.1), we can infer that $g(1) = \frac{1}{50}$, and $g(j) < \frac{1}{n}$ for all $j > 1$.

We now construct f in three steps. Start with $f = g^\theta$.

1. (Preserving the minor) First, set the value of entries of f that are of the form $(x_1; x_2; \dots; x_n)$ as

$$f(x_1; x_2; \dots; x_n) = g(x_1; x_2; \dots; x_n) \quad \forall \mathbf{x} \in \{0,1\}^n$$

We then extend it both upwards and downwards i.e. if $f(\mathbf{p})$ is set to 1 and $\mathbf{p} \geq \mathbf{q}$, then set $f(\mathbf{q}) = 1$ as well, and similarly, if $f(\mathbf{q})$ is set to 0, and $\mathbf{p} \geq \mathbf{q}$, then we set $f(\mathbf{p}) = 0$. This ensures that g is a minor of f and that the partial function f is monotone.

2. (Destroying the influence of 1) Next, we ensure that the Shapley value of the coordinate 1 is low by the following operation: consider all \mathbf{y} such that $f(\mathbf{y})$ has not been set in the first step, $y_1 = 0$ and $f(1; y_2; \dots; y_{2n-1})$ is already set to 1 in the first step. Then set $f(\mathbf{y})$ to be 1. Similarly, if \mathbf{y} satisfies $y_1 = 1$ and $f(0; y_2; \dots; y_{2n-1})$ is already set to 0 in the first step, set $f(\mathbf{y})$ to be 0 if it has not been set in the first step.

We claim that the updated partial function f is still a monotone partial function. Consider $\mathbf{p}; \mathbf{q} \in \{0,1\}^{2n-1}$ such that $\mathbf{p} \geq \mathbf{q}$. Suppose that $f(\mathbf{p})$ is set to be 1. If it is set in the first step, as we extended the partial function upwards in the first step, $f(\mathbf{q}) = 1$ as well. If

$f(\mathbf{p})$ is set to be 1 in the second step, it implies that $f(\mathbf{p}^\theta)$ has been set to 1 in the first step, where \mathbf{p}^θ is obtained from \mathbf{p} by setting p_1 to be 1. Let $\mathbf{q}^\theta \geq f(0; 1)g^{2n-1}$ be obtained from \mathbf{q} by setting $q_1 = 1$. As $\mathbf{p}^\theta \leq \mathbf{q}^\theta$, $f(\mathbf{q}^\theta)$ has been set to 1 in the first step as well. Thus, $f(\mathbf{q})$ is set to be 1 in the second step. The same argument can be used to show that if $f(\mathbf{q}) = 0$, then $f(\mathbf{p}) = 0$ as well.

3. (Adding influence to 2) For all \mathbf{y} for which $f(\mathbf{y}) = 0$ set $f(\mathbf{y}) = y_2$. The fact that the final function f is monotone follows from observing that any completion of a partial monotone function using a monotone function results in a monotone function.

Finally, our goal is to argue about the Shapley value of the coordinates of the function f . First, we show that the Shapley value of the coordinate 1 in f is $o(1)$. Suppose there exists $\mathbf{p} = (0; y_2; y_3; \dots; y_{2n-1})$ and $\mathbf{q} = (1; y_2; y_3; \dots; y_{2n-1})$ such that $f(\mathbf{p}) = 0$ and $f(\mathbf{q}) = 1$. We claim that both the values $f(\mathbf{p})$ and $f(\mathbf{q})$ are set in the first step of the above procedure. Suppose for contradiction that this is not the case. If neither of them is set in the first step, then they will not be set in the second step either, and in the third step, both of them will be assigned the same value, a contradiction. If exactly one of them is set in the first step, then in the second step, the other value would be set to be equal to it, a contradiction as well. Thus, both the values $f(\mathbf{p})$ and $f(\mathbf{q})$ are set in the first step.

Let $B = B_g(1) = f(0; 1)g^{n-1}$ be the boundary of the coordinate 1 in g . As $f(\mathbf{q})$ is set to be 1 in the first step, there exists $\mathbf{x} \geq f(0; 1)g^n$ such that $g(\mathbf{x}) = 1$ and $(x_1; x_2; x_3; \dots; x_n) \leq \mathbf{q}$. As \mathbf{x} is not less than or equal to \mathbf{p} , we can conclude that $x_1 = 1$ and $g(0; x_2; x_3; \dots; x_n) = 0$. In other words, $(x_2; x_3; \dots; x_n) \geq B$. Similarly, there exists \mathbf{x}^θ such that $g(\mathbf{x}^\theta) = 0$ and $(x_1^\theta; x_2^\theta; x_3^\theta; \dots; x_n^\theta) \leq \mathbf{p}$. By the same argument as above, we can conclude that $(x_2^\theta; x_3^\theta; \dots; x_n^\theta) \geq B$. Combining the both, we can conclude that there exist $\mathbf{x}; \mathbf{x}^\theta \geq B$ such that $(x_2; x_2; x_3; x_3; \dots; x_n; x_n) \leq (y_2; y_3; \dots; y_{2n-2}) \leq (x_2^\theta; x_2^\theta; x_3^\theta; x_3^\theta; \dots; x_n^\theta; x_n^\theta)$. Note that if the above inequality is true for a $(y_2; y_3; \dots; y_{2n-2})$, we directly get that $(y_2; y_3; \dots; y_{2n-2})$ is in the boundary of the coordinate 1 in f .

Observe that the boundary of coordinate 1 in g is the set of vectors $(x_2; x_3; \dots; x_n)$ such that $\sum_{j=2}^n x_j \geq \frac{51}{100}n$. By the previous argument, we can deduce that the boundary $B^f = B_f(1)$ of the coordinate 1 in f is the set of vectors $\mathbf{y} = (y_2; y_3; \dots; y_{2n-1})$ that satisfy the following property: The number of $i \in [n-1]$ such that both $y_{2i} = y_{2i+1} = 1$ is at least $\frac{49}{100}n$. Similarly, the number of $i \in [n-1]$ such that $y_{2i} = y_{2i+1} = 0$ is at least $\frac{49}{100}n$. Observe that this implies that we require that $\sum_{j=2}^{2n-1} y_j \geq \frac{51}{50}n$. However, for every integer l such that $\frac{49}{50}n \leq l \leq \frac{51}{50}n$, when we sample a uniformly random vector $\mathbf{y} = (y_2; y_3; \dots; y_{2n-1})$ with $\sum_{j=2}^{2n-1} y_j = l$, the probability that the number of $i \in [n-1]$ such that both $y_{2i} = y_{2i+1} = 1$ is at least $\frac{49}{100}n$ is $o(\frac{1}{n})$. Thus, using Equation (3.1), we can infer that the Shapley value of the coordinate 1 in f is $o(1)$.

We now show that the coordinate 2 has $o(1)$ Shapley value in f . Consider $\mathbf{y} = (y_1; y_2; \dots; y_{2n-1})$ such that $\frac{49n}{50} < \text{hw}(\mathbf{y}) \leq \frac{51n}{50}$. If the number of i such that both $y_{2i} = y_{2i+1} = 1$ is less than $\frac{49}{100}n$, we have $(y_1; y_3; \dots; y_{2n-1}) \geq B_f(2)$. However, for every integer l such that $\frac{49}{50}n \leq l \leq \frac{51}{50}n$, when we sample a uniformly random \mathbf{y} with $\text{hw}(\mathbf{y}) = l$, with probability $1 - o(1)$, the number of i such that both $y_{2i} = y_{2i+1} = 1$ is less than $\frac{49}{100}n$. Thus, using Equation (3.1), we can infer that the Shapley value of the coordinate 2 is $o(1)$ in the function f . Finally, by symmetry, we can observe

that $f(i) = f(3)$ for all $i \geq 3$, and thus, $f(i) = o(1)$ for all $i \geq 3$.

□

Chapter 4

d -to-1 Hardness of Coloring 3-colorable graphs with $O(1)$ colors

4.1 Introduction

Determining if a graph is 3-colorable is one of the classic NP-complete problems. Thus, given a 3-colorable graph it is NP-hard to color it with 3 colors. The best known polynomial time algorithms for coloring 3-colorable graphs use about $n^{0.2}$ colors, where n is the number of vertices in the graph [KT17]. On the other hand, on the hardness front, we only know that 5-coloring 3-colorable graphs is NP-hard [Bar+21].

This embarrassingly large gap between the hardness and algorithmic results has prompted the quest for conditional hardness results for approximate graph coloring. The canonical starting point for most strong inapproximability results is the Label Cover problem. The Unique Games Conjecture of Khot [Kho02a], which asserts strong inapproximability of Label Cover when the constraint maps are bijections, has formed the basis of numerous tight hardness results for problems which have defied NP-hardness proofs. However, the imperfect completeness inherent in the Unique Games Conjecture makes it unsuitable as the basis for hardness results for graph coloring, where we want *all* edges to be properly colored under the coloring.

In [Kho02a], along with the Unique Games Conjecture, Khot introduced the d -to-1 conjecture. The d -to-1 conjecture says that given a Label Cover instance whose constraint relations are d -to-1 functions, it is NP-hard to decide if there exists a labelling that satisfies all the constraints or no labelling can satisfy even an ϵ fraction of constraints, for arbitrarily small $\epsilon > 0$. (The key is that d can be held fixed and achieve soundness $\epsilon \neq 0$.) Constraints similar to 2-to-1 also played an implicit role in the beautiful work of Dinur and Safra on inapproximability of vertex cover [DS05].

Based on the 2-to-1 conjecture, Dinur, Mossel and Regev [DMR09], extending the invariance principle based techniques of [Kho+07; MOO10], proved the hardness of coloring graphs that are promised to be 4-colorable with any constant number of colors. Furthermore, they prove the same for 3-colorable graphs under a certain “fish shaped” variant of the 2-to-1 conjecture. In this work,

we prove that the same result can be proved under the weaker assumption of d -to-1 conjecture¹, for some (arbitrarily large) constant d .

Theorem 26. *Assume that d -to-1 conjecture is true for some constant d . Then, for every positive integer $t \geq 3$, it is NP-hard to color a 3-colorable graph G with t colors.*

We stress that the d -to-1 conjecture insists on perfect completeness (i.e., hardness on satisfiable instances), and this important feature seems necessary for its implications for coloring problems, where we seek to properly color all edges. The variant of the 2-to-1 conjecture where one settles for near-perfect completeness was recently established in a striking sequence of works [KMS17; Din+18a; Din+18b; KMS18].

The result of [DMR09] in fact shows hardness of finding an independent set of density ϵ in a 3-colorable graph for arbitrary $\epsilon > 0$ (which immediately implies the hardness of finding a coloring with $1 - \epsilon$ colors). Our result in Theorem 26 above does *not* get this stronger hardness for finding independent sets. But it is conditioned on the d -to-1 conjecture for arbitrary d rather than the specific 2-to-1 conjecture. We note that proving the d -to-1 conjecture for some large d could be significantly easier than the 2-to-1 conjecture, so Theorem 26 perhaps provides an avenue for resolving a longstanding challenge concerning the complexity of approximate graph coloring.

Our proof of Theorem 26 is a simple combination of two results. First, following the methodology of [DMR09], we prove that the d -to-1 conjecture implies that coloring a $2d$ -colorable graph with $O(1)$ colors is NP-hard. The result of [DMR09] is the $d = 2$ case of this claim. In fact, they state in a future work section that the d -to-1 conjecture should imply hardness of $O(1)$ -coloring q -colorable graphs for some large enough $q = q(d)$. However, they did not specify the details of the reduction or an explicit value of q , and mention determining the dependence of q on d as an interesting question. Here we show the conditional hardness based on d -to-1 conjecture holds for $q = 2d$ (achieving $q < 2d$ seems unlikely with the general reduction approach of [DMR09]).

The key technical ingredient necessary for such a reduction is a symmetric Markov chain on $[q]^d$ where transitions are allowed only between disjoint tuples and which has spectral radius bounded away from 1. We show the existence of such a symmetric Markov chain for $q = 2d$. We do so via a connection to matrix scaling, which enables us to deduce the necessary chain at a conceptual level without messy calculations. Specifically, we use the result [CD72], which follows from the Sinkhorn-Knopp iterative matrix scaling algorithm [SK67], that if a non-negative symmetric matrix A has *total support* then there is a symmetric doubly stochastic matrix supported on the non-zero entries of A . When A is the adjacency matrix of a graph G , the total support condition is equivalent to every edge of G belonging to a cycle cover. We describe a graph on $[q]^d$ whose edges connect disjoint tuples and where every edge belongs to a cycle cover.

Our second ingredient is a remarkable yet simple reduction due to Krokhn, Opršal, Wrochna and Živný [Kro+20], which exploits the relation between the arc-chromatic number and chromatic number of a digraph [PR81]. Let $b : \mathbb{N} \rightarrow \mathbb{N}$ be defined by $b(n) := \frac{n}{bn=2c}$. Their result then is that $b(t)$ -coloring $b(c)$ -colorable graphs is polynomial time (in fact logspace) reducible to

¹For d -to-1 Label Cover, there are two definitions possible, one where the constraint maps are *at most* d -to-1 with each element in the range having at most d pre-images, and one where the constraint maps are *exactly* d -to-1. In this chapter, we stick with the exact variant.

t -coloring c -colorable graphs. Since $b(n)$ is increasing and $b(n) > n$ for all $n \geq 4$, it follows that a NP-hardness result for $O(1)$ -coloring q -colorable graphs also implies NP-hardness of $O(1)$ -coloring 4-colorable graphs. Furthermore, the NP hardness of $O(1)$ -coloring of 3-colorable graphs follows from the above by applying the arc graph reduction twice to K_4 .

Overview. In Section 4.2, we define the Label Cover problem, and state the d -to-1 conjecture formally. We also introduce low degree influences that we need later. In Section 4.3, we first prove the existence of the Markov chain with required properties, and then describe the reduction from Label Cover to Approximate Coloring. We note that the reduction is in fact exactly the same one used in [DMR09], the difference being in using a different Markov Chain. For the sake of completeness, we present the reduction and the preliminaries required.

4.2 Preliminaries

We first formally define the hardness conjectures.

4.2.1 d -to-1 Conjecture

We first state the d -to-1 conjecture. As is the case with [DMR09], we will state and use the *exact* d -to-1 variant where the constraint maps have exactly d pre-images for each element in the range. Khot's original formulation only required that there are at most d pre-images for each element in the range. The d -to-1 conjecture becomes stronger for smaller d (so that the 2-to-1 is the strongest form of the conjecture)—this is obvious for the variant where the maps are at most d -to-1. For the exact variant, if we allow the Label cover graph to have multiple edges, we can reduce d -to-1 conjecture to $(d + 1)$ -to-1 conjecture using a simple argument. We present this reduction in Section 4.4. On that note, we remark without details that our reduction indeed works with the multigraph variant of d -to-1 conjecture.

Conjecture 27. (*(Exact) d -to-1 Conjecture*) For every $\epsilon > 0$, given a bipartite Label Cover instance $G = ((V = X \sqcup Y; E); (dR; R); \rho)$ satisfying the following constraints:

- (i) We refer to X as the vertices on the left, and Y as the set of vertices on the right. The vertices belonging to X are to be assigned labels from $[dR]$ while the vertices in Y are to be assigned labels from $[R]$.
- (ii) The constraints are d -to-1 i.e. for every $b \in [R]$, there are precisely d values $a \in [dR]$ such that $(a; b) \in \rho_e$ for every relation ρ_e in the instance.

It is NP-hard to distinguish between the following cases:

1. There is a labeling that satisfies all the constraints in G .
2. No labeling can satisfy more than ϵ fraction of constraints in G .

Similar to the d -to-1 constraints, one can consider d -to- d constraints in the Label Cover. In order to do so, we define the relation $d \text{ \$ } d$ on $[dR] \times [dR]$:

$$d \text{ \$ } d = \{(d_i, p+1; d_i, q+1) \mid i \in [R]; 1 \leq p, q \leq dg\}$$

A constraint $C = (a; b) \in [dR] \times [dR]$ is said to be *d-to-d* if there exist permutations π_1 and π_2 on $[dR]$ such that $(a; b) \in C$ iff $(\pi_1^{-1}(a); \pi_2^{-1}(b)) \in C$.

In [DMR09], it is proved that Conjecture 27 implies the following conjecture.

Conjecture 28. (d-to-d conjecture) For every $\epsilon > 0$ and every $t \in \mathbb{N}$, there exists $d \in \mathbb{N}$ such that given a Label Cover instance $C = ((V; E); dR; \cdot)$ where all the constraints are *d-to-d*, it is NP-hard to distinguish between the following cases:

- (i) $\text{sat}(G) = 1$, or
- (ii) $\text{isat}_t(G) < \epsilon$

Here, $\text{sat}(G)$ denotes the maximum fraction of constraints satisfied by any labeling. Similarly, $\text{isat}(G)$ denotes the size of the largest set $S \subseteq V$ such that there exists a labeling that satisfies all the constraints induced on S . The value $\text{isat}_t(G)$ denotes the size of largest set $S \subseteq V$ such that there exists a labeling that assigns at most t labels to each vertex that satisfies all the constraints induced on S . A constraint between $u; v$ is said to be satisfied by labeling assigning multiple labels to u and v if and only if there exists at least one pair of labels s and t among the multiple labels that satisfy the constraint.

4.2.2 Low degree in uences

Next, we define the low degree in uences that we need later. We refer the reader to [DMR09] for a comprehensive treatment of the same.

Let $e_0 = 1; e_1; \dots; e_{q-1}$ be an orthonormal basis for \mathbb{F}_q^n . We can define the set of functions $f_x : [q]^n \rightarrow \mathbb{R}; x \in [q]^n$ as $f_x(y) = (e_{x_1}(y_1); e_{x_2}(y_2); \dots; e_{x_n}(y_n))$. Observe that these functions form a basis for the functions from $[q]^n$ to \mathbb{R} . Let $\langle f; g \rangle = \sum_{x \in [q]^n} f(x)g(x)$, where we define the inner product between functions $f; g : [q]^n \rightarrow \mathbb{R}$ as $\langle f; g \rangle = \sum_{x \in [q]^n} f(x)g(x)$. We define the low degree in uence of f as follows:

Definition 29. For a function $f : [q]^n \rightarrow \mathbb{R}$, the *low degree in uence* of the coordinate i is defined as follows:

$$I_i^k(f) = \sum_{x: x_i \in \{0; \dots; k\}} f^2(x)$$

Note that the above definition is independent of the basis $e_1; \dots; e_{q-1}$ that we start with, as long as $e_0 = 1$. From the above definition, we can infer that for functions $f : [q]^n \rightarrow [0; 1]$, the sum of low degree in uences is bounded by

$$\sum_i I_i^k(f) \leq k$$

For a vector $x \in [q]^{dR}$, let $\bar{x} \in [q]^{dR}$ be the corresponding element in $[q]^{dR}$ i.e.

$$\bar{x} = ((x_1; x_2; \dots; x_d); (x_{d+1}; x_{d+2}; \dots; x_{2d}); \dots; (x_{dR-d+1}; x_{dR-d+2}; \dots; x_{dR}))$$

Similarly, for $y \in [q]^{dR}$, let \underline{y} denote the inverse of above operation. We can extend this notion to functions as well: For a function $f : [q]^{dR} \rightarrow \mathbb{R}$, let the function $\bar{f} : [q]^{dR} \rightarrow \mathbb{R}$ be defined naturally by

$$\bar{f}(y) = f(\underline{y})$$

Similarly, for a function $f : [q]^d \rightarrow \mathbb{R}$, let $\bar{f} : [q]^{dR} \rightarrow \mathbb{R}$ be defined as $\bar{f}(x) = f(\bar{x})$.

We need the following lemma:

Lemma 30. For any function $f : [q]^d \rightarrow \mathbb{R}$ and any $k \geq 1$ and $i \geq 1$,

$$I_i^k(\bar{f}) = \sum_{j=1}^d I_{d_i}^{dk_{d+j}}(f)$$

Proof. Fix a basis $\{x_j\}$ of functions from $[q]^d \rightarrow \mathbb{R}$ as above. The functions $\{x_j\}$ form a basis for functions from $[q]^d \rightarrow \mathbb{R}$, where $x_j(\bar{y}) = x_j(y)$. Note that $\bar{f}^k(x) = f^k(\bar{x})$. Thus we get

$$\begin{aligned} \sum_i I_i^k(\bar{f}) &= \sum_{x: x_i \in \{0, \dots, 0\}; j \in [k]} \bar{f}^k(x) = \sum_{x: x_i \in \{0, \dots, 0\}; j \in [k]} f^k(\bar{x}) \\ &= \sum_{j=1}^d \sum_{x: x_{d_i} \in \{0, \dots, 0\}; j \in [k]} f^k(\bar{x}) \\ &= \sum_{j=1}^d I_{d_i}^{dk_{d+j}}(f) \quad \square \end{aligned}$$

Using the invariance principle and Borell's inequality, [DMR09] prove the following:
 Theorem 31. Let q be a fixed integer, and T be a symmetric Markov chain on $[q]$ with $r(T) < 1$. Then for every $\epsilon > 0$, there exists a $\delta > 0$ and a positive integer k such that the following holds: For every $f, g : [q]^n \rightarrow [0, 1]$ if $E[f] > \delta$; $E[g] > \delta$ and $f, T^k g = 0$, then

$$|I_i^k(f) - I_i^k(g)| \leq \epsilon$$

where $r(T)$ denotes the second largest eigenvalue (in absolute value) of

4.3 d-to-1 hardness for 3-colorable graphs

In this section, we will prove Theorem 26.

4.3.1 Reducing chromatic number to 3

The following lemma is present in [Kro+20] based on a beautiful result concerning the arc-chromatic numbers of digraphs from [PR81].

Lemma 32. (Theorem 1.8 of [Kro+20]) Suppose there exists N such that $O(1)$ coloring q -colorable graphs is NP-hard. Then $O(1)$ coloring 3-colorable graphs is NP-hard.

Let $\text{Graph-Coloring}(t; c)$ denote the promise problem of distinguishing if a graph can be colored with c colors, or cannot even be colored with c colors. The statement is proved by presenting a reduction from $\text{Graph-Coloring}(b(t); b(c))$ to $\text{Graph-Coloring}(t; c)$ in polynomial time, for the function $b(n) := \lfloor \frac{n}{b} \rfloor$. The reduction works by constructing the arc-graph of the underlying graphs, and using the property of arc graphs that the chromatic number of the arc graph can be bounded precisely using the chromatic number of the original graph. Since b is an increasing function and $b(n) > n$ for all $n \geq 4$, setting $c = 4$ and t large enough proves the statement claimed in the lemma. The reduction from b -colorable graphs to c -colorable graphs is achieved by applying the arc graph construction twice recursively.

Thanks to Lemma 32, we can restrict ourselves to the weaker goal of proving that coloring q -colorable graphs is NP-hard for some fixed constant q , assuming Conjecture 27. In fact, following [DMR09], we prove a stronger statement showing hardness of finding independent sets of fraction ϵ of vertices for any $\epsilon > 0$. Combined with Lemma 32, this immediately gives us Theorem 26.

Theorem 33. Suppose that Conjecture 28 is true for a constant d . Then, there exists a constant $q = q(d)$ such that for every $\epsilon > 0$, given a graph G , it is NP-hard to distinguish the following cases:

1. G can be colored with q colors.
2. G does not have any independent set of relative size ϵ .

In fact, we can take $q = 2d$.

In the remainder of the section, we will prove Theorem 33. We next develop the main technical ingredient that we will plug into the reduction framework of [DMR09] to establish Theorem 33.

4.3.2 A symmetric Markov chain supported on disjoint tuples

A Markov chain T defined on a state space Σ is said to be symmetric if the transition matrix T is symmetric, namely for all pairs of states $x, y \in \Sigma$, the probability of transition from x to y is equal to the probability of transition from y to x . Symmetry of the Markov chain ensures that the uniform distribution is stationary which is essential when we compose the Label Cover-Long Code reduction with the Markov chain. We define the spectral radius $r(T)$ of a symmetric Markov chain as the second largest eigenvalue in absolute value of its transition probability matrix, i.e., if $1 = \lambda_1 \geq \lambda_2 \geq \dots \geq \lambda_q$ are the eigenvalues, then $r(T) = \max\{|\lambda_j|; j \geq 2\}$.

We now show the existence of a symmetric Markov Chain T on $[q]^d$ with $r(T) < 1$ if $d \geq 2; q \geq 2d$. Furthermore, there is a nonzero transition probability between two elements $x, y \in [q]^d$ only if the support of x and y are disjoint. In [DMR09], such a Markov Chain is shown to exist for the values $(q; d) = (3; 1); (4; 2)$.

Lemma 34. Suppose that $d \geq 2; q \geq 2d; d \geq 2$. There exists a symmetric Markov chain T on $[q]^d$ such that $r(T) < 1$. Furthermore, if the transition matrix T has positive probability in T , then $\{x_1, x_2, \dots, x_d\} \cap \{y_1, y_2, \dots, y_d\} = \emptyset$.

Proof. We first construct an undirected graph G on $[q]^d$ such that there is an edge between $x, y \in [q]^d$ only if the support of x and y are disjoint. We then use a matrix scaling algorithm to

obtain a symmetric Markov chain from the adjacency matrix A . For the resulting Markov chain to have $\alpha(T) < 1$, we need that the underlying graph is connected, and is not bipartite. Furthermore, for the scaling algorithm to produce a valid Markov chain, we need that every edge of G is present in a cycle cover, where a cycle cover of a graph is a disjoint union of cycles that covers every vertex in the graph. Note that we allow trivial 2-cycles in a cycle cover, where we just take an edge twice.

We say that two multisets $x = (x_1; x_2; \dots; x_d); y = (y_1; y_2; \dots; y_d) \in [q]^d$ are of the same type if the following condition holds: for all pairs of indices $j \in [d]$, $x_i = x_j$ if and only if $y_i = y_j$ and $(x_i - x_j)(y_i - y_j) = 0$. Note that this is an equivalence relation, and thus each element $x \in [q]^d$ uniquely determines its type.

Consider the graph $G = (V; E)$ where the vertex set $V = [q]^d$. We add two kinds of edges in this graph. We add an edge between every pair $x, y \in [q]^d$ that are of the same type, and have disjoint support. Let the subset S of elements that are supported on single element be denoted by S , i.e.,

$$S = \{(1; 1; \dots; 1); (2; 2; \dots; 2); \dots; (q; q; \dots; q)\}$$

We also add edges between x and y if their support is disjoint, and at least one of x and y belongs to S .

First, we claim that G is connected. This follows from the fact that the set of nodes S are connected to each other, and every vertex x is adjacent to at least one vertex $s \in S$. As $q \geq 4$, the graph is not bipartite (indeed S induces a q -clique). We will now prove that every edge in this graph is part of a cycle cover. Given an undirected graph on vertex set V , a cycle cover of it is a function $\sigma : V \rightarrow V$ that is bijective, and $(u, v) \in E$ only when u and v are adjacent in the underlying graph.

Towards this, we first prove that for every edge G between multisets of the same type, there is a cycle cover that uses that edge. For each type, consider the graph obtained by taking the vertices as multisets of that type, and with edges between two multisets of the same type if they are disjoint. Note that for every type, this graph is isomorphic to a Kneser graph $K(q, k)$ (for some $k \leq d$), whose vertex set corresponds to element subsets of $[q]$ and there is an edge between two subsets if they are disjoint.

By symmetry across the subsets, we can infer that the Kneser graphs are regular. Note that every regular graph contains a cycle cover: For a regular graph H , consider a bipartite graph H^0 which contains a copy of H on both the left side L , and right side R . There is an edge between $x \in L; y \in R$ of H^0 if and only if x, y are adjacent in H . As H is a regular graph, H^0 is a regular bipartite graph, and thus, contains a perfect matching. This perfect matching directly gives a cycle cover of H . Furthermore, as Kneser graphs are also vertex-transitive, every edge in these graphs is part of a cycle cover.

Next, we consider edges G that are between multisets of different types i.e. edges between multisets $x; y$ where exactly one of x and y is in S . Consider an edge between $s \in S$ and $x \in V \setminus S$. As $q \geq 2d$, every multiset in G is adjacent to at least one multiset of the same type y but a multiset that is adjacent to x in G and is of the same type as x . Let $s^0 \in S$ be chosen such that

it is adjacent to y in G . As S is a complete subgraph of G , s and s^0 are adjacent in G . From the previous argument about edges between multisets of the same type, we can infer that there is a cycle cover of G where y is mapped to x , and s is mapped to s^0 . We can modify this cycle cover by transforming it as follows: $(s \rightarrow x)$ can be made part of cycle cover by transforming $(s \rightarrow s^0); (y \rightarrow x)$ to $(s \rightarrow x); (y \rightarrow s^0)$ and keeping rest of the cycle cover intact. Thus, we have proved that every edge of G is part of a cycle cover.

Let A denote the adjacency matrix of the above graph. Using the Sinkhorn Knopp iterative algorithm, it is proved in [CD72] that if a non-negative symmetric matrix A has total support, then there exists a diagonal matrix D such that DAD is a doubly stochastic matrix. A square matrix $A = (a_{ij})$ of order n is said to have total support if $a_{ij} > 0$, and for every nonzero entry a_{ij} of A , there exists a permutation of $[n]$ such that $(i) = j$ and for all $e \in [n]$; $a_{e; (e)} > 0$. When the matrix A is an adjacency matrix of a graph, the total support condition translates to the requirement that every edge of G is part of a cycle cover, a property we have already shown to hold for the graph G .

Thus, we can apply the above scaling result, and view the resulting matrix DAD as the transition matrix of a Markov chain. As A and D are symmetric, T is symmetric. As A is connected and no principal diagonal element of D is zero, T is connected as well. Note that every nonzero element of A stays nonzero in T , and A is not bipartite. The above two facts combined ensure that the spectral radius of T is strictly less than 1. We conclude that there exists a symmetric Markov chain on state space $[q]^d$ that has both the properties: (i) $r(T) < 1$, and (ii) there is nonzero probability of transition between two multisets only when their support is disjoint. \square

4.3.3 Proof of Theorem 33

Let d be the constant for which Conjecture 27 is true. Thus, Conjecture 28 is true for the same d valued as well. Choose ϵ ; T from Lemma 34 such that T is a symmetric Markov chain on $[q]^d$ such that $r(T) < 1$.

We now reduce the given d -Label Cover instance to the problem of finding independent sets in q -colorable graphs. To be precise, given a Label Cover instance $(V; E; d; R; \rho)$, we output a graph $G^0 = (V^0; E^0)$ such that

1. Completeness: If $\text{val}_\rho(G) \geq \epsilon$, G^0 can be colored with q colors.
2. Soundness: If $\text{val}_\rho(G) < \epsilon$, then G^0 does not have any independent set of size $\Omega(d)$.

The parameters ϵ and $\Omega(d)$ will be set later.

Reduction. Our reduction follows the standard Label Cover Long Code paradigm, and in particular closely mirrors [DMR09]. We replace each vertex $v \in V$ of the Label Cover with a set f_w of $[q]^{dR}$ nodes, each corresponding to a vertex w in G . Consider an edge $e = (u; v)$ where ρ_e is an associated constraint with permutations ρ_1, ρ_2 on $[dR]$ such that $(a; b) \in \rho_e$ if and only if $(\rho_1^{-1}(a); \rho_2^{-1}(b)) \in \rho_e$.

We add an edge between $(x_1, x_2, \dots, x_{dR}) \in f_u$ and $(y_1, y_2, \dots, y_{dR}) \in f_v$ to E^0 if and only if

$$\sum_{i \in [R]} T((x_{1(d_i-d+1)}; x_{1(d_i-d+2)}; \dots; x_{1(d_i)}) \neq (y_{2(d_i-d+1)}; y_{2(d_i-d+2)}; \dots; y_{2(d_i)})) > 0$$

Completeness. Suppose $\sigma : V \rightarrow [q]$ be a labeling satisfying all the constraints of the Label Cover instance \mathcal{C} . We color the nodes $(x_1, x_2, \dots, x_{dR}) \in f_w$ with $x_{(w)} \in [q]$. We claim that this is a legit q -coloring of G^0 . Suppose that we added an edge between f_u and f_v . Let x be colored with x_a and y be colored with y_b . As $(a, b) \in (u, v)$, we have $(\sigma^{-1}(a), \sigma^{-1}(b)) \in d \in \mathcal{D}$. Thus, there exist $i \in [R]; 1 \leq p, q \leq d$ such that $a = x_{1(d_i-d+p)}$ and $b = y_{2(d_i-d+q)}$. As we have added an edge between f_u and f_v , $x_a \neq y_b$ as the Markov chain has nonzero probability only between two elements $\sigma^{-1}(a)$ with disjoint support. Thus, there exists a q -coloring of G^0 when \mathcal{C} is satisfiable.

Soundness. We prove the contrapositive that if G^0 has an independent set of relative size ϵ then there exists a labeling σ with $\text{sat}_t(\mathcal{C}) \geq \epsilon$. Let $S \subseteq V^0$ be the largest independent set of G^0 . We know that $|S| \geq \epsilon |V^0|$. This implies that in at least ϵ fraction of the long code blocks, at least $\frac{\epsilon}{2}$ fraction of nodes belong to S . Let this subset of V be denoted by Z . Our goal is to show that there exists a small set of labels $Z \subseteq [q]^{dR}$ to which we can decode the vertices in Z such that all the constraints induced in Z are satisfied by σ .

For every vertex $w \in Z$, we define functions $g_w : [q]^{dR} \rightarrow \{0, 1\}$ to be the indicator functions of S inside the long code blocks corresponding to w . $g_w(x) = 1$ if and only if $x \in S$. Consider an edge $e = (u, v)$ corresponding to the constraint t_e induced in Z . Let the functions $f : [q]^{dR} \rightarrow \{0, 1\}$ and $g : [q]^{dR} \rightarrow \{0, 1\}$ be defined such that $f(x^{-1}) = g_u(x)$ and $g(y^{-2}) = g_v(y)$, where σ_1 and σ_2 are the permutations underlying the relation $i.e. (a, b) \in e$ if and only if $(\sigma_1^{-1}(a), \sigma_2^{-1}(b)) \in d \in \mathcal{D}$.

We note that $\sum_{i \in [R]} T(g_i)$ is equal to zero. In other words, suppose that $x \in [q]^{dR}; x \in f_u; y \in f_v$ are such that

$$\sum_{i \in [R]} T((x_{1(d_i-d+1)}; x_{1(d_i-d+2)}; \dots; x_{1(d_i)}) \neq (y_{2(d_i-d+1)}; y_{2(d_i-d+2)}; \dots; y_{2(d_i)})) > 0: \quad (4.1)$$

Then, $f(x)g(y) = 0$. Suppose for contradiction that there exist $x \in [q]^{dR}$ satisfying the above condition, and $f(x) = g(y) = 1$. Let $x^0 \in f_u; y^0 \in f_v$ be such that $(x^0)^{-1} = x; (y^0)^{-2} = y$. We have $g_u(x^0) = g_v(y^0) = 1$. That is, both $x^0 \in f_u; y^0 \in f_v$ are in the independent set S . However, Equation (4.1) can be rewritten as the following:

$$\sum_{i \in [R]} T((x_{1(d_i-d+1)}^0; x_{1(d_i-d+2)}^0; \dots; x_{1(d_i)}^0) \neq (y_{2(d_i-d+1)}^0; y_{2(d_i-d+2)}^0; \dots; y_{2(d_i)}^0)) > 0: \quad (4.2)$$

Note that this is precisely the condition for adding edges in G^0 . Thus, Equation (4.2) implies that $x^0 \in f_u$ and $y^0 \in f_v$ are adjacent in E^0 , and thus cannot both be part of the independent set S . This completes the proof that $\sum_{i \in [R]} T(g_i) = 0$.

Thus, $\sum_{i \in [R]} T(g_i)$ is also equal to zero, where $\bar{f} : [q]^{dR} \rightarrow \{0, 1\}$ and $\bar{g} : [q]^{dR} \rightarrow \{0, 1\}$ are the corresponding functions $\sum_{i \in [R]} T(\bar{f}_i)$ of $f; g$. From the definition of Z , $E(\bar{f}) \geq \frac{\epsilon}{2}$ and $E(\bar{g}) \geq \frac{\epsilon}{2}$. We apply Theorem 31 to \bar{f} and \bar{g} to deduce that there exist $k \in [R]$, a positive integer $k = k(\epsilon)$ and

$= (i_1, i_2)$ such that $i_1 \in [dR]$ and $i_2 \in [dR]$. This motivates us to define the label set of vertex $w \in Z$, $L(w)$ as the following -

$$L(w) := \{i \in [dR] : I_i^{dk}(g_w) \neq \bar{g}\}$$

As the sum of degree in edges of all variables is at most kd , the size of $L(w)$ is upper bounded by kd for every w . Thus, we set the parameter ϵ to be $\frac{1}{kd}$.

The final step is to prove that the labeling is indeed a valid labeling inside edges induced in Z . Consider an edge $e = (u; v)$ induced in Z with the constraint relation being e such that $(a; b) \in e$ if and only if $(i_1(a); i_2(b)) \in d \in \mathcal{R}$. Our goal is to show that there exist indices $i_1; i_2 \in [dR]$ such that $i_1 \in L(u)$; $i_2 \in L(v)$ and $(i_1; i_2) \in e$. Using Theorem 31, we can deduce that there exists $i_1 \in [dR]$ such that $i_1 \in L(u)$ and $I_{i_1}^{dk}(f) \neq \bar{g}$. Using Lemma 30, we can conclude that there exists $i_2 \in [dR]$ such that $i_2 \in L(v)$ and $I_{i_2}^{dk}(g) \neq \bar{g}$ such that $(i_1; i_2) \in d \in \mathcal{R}$. Let $i_1; i_2 \in [dR]$ be such that $i_1 \in L(u)$; $i_2 \in L(v)$. As $f(x) = g_u(x)$, $I_{i_1}^{dk}(f) \neq \bar{g}$. And thus, $i_1 \in L(u)$, and similarly $i_2 \in L(v)$. As $(i_1; i_2) \in d \in \mathcal{R}$, $(i_1; i_2) \in e$, which completes the proof.

4.4 Reducing multigraph (exact)d-to-1 to (d + 1)-to-1 conjecture

For the version of d-to-1 conjecture where we only require the constraint maps to be at most, the d-to-1 conjecture trivially implies the (d + 1)-to-1 conjecture. O'Donnell and Wu [OW09] remark that no such reduction appears to be known for the exact conjecture. Here we prove that the exact d-to-1 conjecture implies the exact (d + 1)-to-1 conjecture when the underlying Label Cover instances are allowed to have parallel edges. We remark that multigraph version of exact d-to-1 conjecture, which is implied by the simple graph version, also suffices for our reduction to graph coloring (and indeed all known reductions from Label Cover).

Let $G = ((V = X \cup Y; E); (dR; R); \epsilon)$ be a Label Cover instance such that every constraint is of d-to-1 structure. We reduce it to $G^0 = ((V = X \cup Y; E^0); ((d + 1)R; R); \epsilon^0)$ such that

1. If G is satisfiable, G^0 is satisfiable as well.
2. If every labeling violates at least fraction ϵ of constraints in G , then every labeling violates at least $\epsilon^0 = \frac{\epsilon}{2}$ fraction of constraints in G^0 .

Reduction. We first change the label set \mathcal{X} from $[dR]$ to $[(d + 1)R]$. For every constraint in G between nodes $u \in X$ and $v \in Y$, we replace it with R constraints $i_1; i_2; \dots; i_R$ between u and v in the following way: the relation between old labels is the same as when $x \in dR$, $(x; y) \in i_j$ for $j = 1; 2; \dots; R$ if and only if $(x; y) \in e$. When $x > dR$, $(x; y) \in i_j$ if and only if R divides $(x + j - y)$. This ensures that each new label is mapped to a different label in each of the R new constraints. The constraints are clearly $(d + 1)$ -to-1 form.

Completeness. If there is a labeling satisfying all the constraints of the same labeling satisfies all the constraints in G^0 as well.

Soundness. Suppose that there is no labeling satisfying at least a fraction of constraints in G . Note that this implies that R is at least $\frac{1}{2}$ as there is always a labeling satisfying at least a fraction of constraints: fix a labeling to the vertices on the left, and assign a label to the vertices on the right uniformly at random from $[R]$. We claim that there is no labeling satisfying more than a fraction of constraints in G^0 . Consider an arbitrary labeling $\phi: V \rightarrow [R]$. We can divide the set of edges E^0 of G^0 into two parts: the edges (u, v) such that $(\phi(u) - \phi(v)) \leq dR$ and the edges (u, v) such that $(\phi(u) - \phi(v)) > dR$. Let the set of first type of edges where the left vertex is assigned the new label be denoted by E_1 , and the set of second type of edges be denoted by E_2 . In E_1 , the fraction of constraints that can be satisfied by ϕ is at most $\frac{1}{R}$. Note that we can get a labeling ϕ' of G by replacing labels of vertices x with label greater than dR with an arbitrary label in $[dR]$, and keeping rest of the labels intact. For the edges in E_2 , the labelings ϕ and ϕ' coincide. As ϕ' can satisfy at most fraction of constraints in G , ϕ can only satisfy at most fraction of overall edges in E^0 . Thus, overall ϕ satisfies at most $\frac{1}{R} + \frac{1}{2}$ fraction of constraints in E^0 , which proves the required soundness claim.

Chapter 5

Rainbow coloring hardness via low sensitivity polymorphisms

5.1 Introduction

Similar to the Graph coloring problem, hypergraph coloring also received a lot of attention in Graph Theory and Theoretical Computer Science. Even though there is a simple algorithm to check if a given graph is 2-colorable or not, checking if a k -uniform hypergraph can be colored with two colors so that no hyperedge is monochromatic is one of the classic NP-hard problems. This raises the question of identifying if coloring is easy on special hypergraphs of interest. For example, if a k -uniform hypergraph is k -partite, i.e., the vertices can be partitioned into k parts so that every hyperedge intersects each part, then there are simple algorithms to properly color the hypergraph with two colors. Suppose we know that a k -uniform hypergraph is promised to be $(k-1)$ -partite, can we color it with two colors?

An equivalent way to formulate this question is in terms of rainbow coloring. A k -uniform hypergraph is said to be rainbow colorable if there is a coloring of vertices with colors such that all the colors appear in every edge. Unlike usual coloring, rainbow coloring becomes harder as we have more colors. Note that k -partiteness is the same thing as rainbow colorability. As mentioned above, a k -uniform hypergraph that is promised to be rainbow colorable can be efficiently colored with two colors. One big hammer approach for this is to use semidefinite programming and find a unit vector for each vertex such that sum of the vectors in each edge sum to zero, and then use random hyperplane rounding. But coloring can also be performed by a simple random walk algorithm — start with an arbitrary coloring, and as long as there is a monochromatic edge, pick an arbitrary one and flip the color of a random vertex in it. This process will converge to 2-coloring in a quadratic number of iterations with high probability [McD93].

If we relax the k -rainbow colorability assumption slightly to that of $(k-1)$ -rainbow colorability, there are no known efficient algorithms for coloring. It is tempting to conjecture that in fact this task is hard (in fact, even if we are allowed c colors for any constant c , this was shown assuming the V Label Cover conjecture in [BG17]). If we relax the rainbow colorability assumption further,

Austrin, Bhangale and Potukuchi proved that it is NP-hard to color a k -uniform hypergraph when it is promised to be $(k-2)$ -rainbow colorable [ABP20]. They also showed that it is NP-hard to 2-color a 4-uniform hypergraph even if it is 3-rainbow colorable. In this work, we focus on hardness results for the $(k-1)$ -rainbow colorable case, as this promise is the closest to k -partiteness which makes coloring easy. While we can't show hardness of coloring, we show that rainbow coloring with $\frac{k-2}{2}e$ colors is hard. Formally, our main result is the following.

Theorem 35. Fix an integer $k \geq 4$. Given a k -uniform hypergraph that is promised to be $(k-1)$ -rainbow colorable, it is NP-hard to rainbow color it with $\frac{k-2}{2}e$ colors.

As a corollary, we also get the following, which extends the similar result of [ABP20] for the $k=4$ case (their techniques did not generalize beyond the 4-uniform case).

Theorem 36. For $k \geq 6$, given a k -uniform hypergraph that is promised to be $(k-1)$ -rainbow colorable, it is NP-hard to 2-color it.

5.1.1 Techniques

There have been broadly three lines of attack on proving hardness for graph and hypergraph coloring problems.

The first line of work gives reductions from Label Cover analyzed using Fourier-analytic techniques of the sort originally pioneered by Hastad [Has01]. Early applications of this method showed strong hardness results for coloring 2-colorable hypergraphs of low uniformity with any constant number of colors [GHS02; Hol02; Kho02a; Sak14]. This approach, augmented with the invariance principle of [MOO10] and some of its extensions such as [DMR09; Mos10; Wen13], was used to prove further hardness results for hypergraph coloring [BK10; GL18] and strong conditional hardness results for graph coloring [DMR09]. These methods usually also prove a stronger statement about finding independent sets in the graphs or hypergraphs. For rainbow coloring, it is proved in [GL18] by combining together many of these techniques that a $k=2$ -rainbow colorable k -uniform hypergraph cannot be colored with any constant number of colors in polynomial time unless $P=NP$. While our results in Chapter 4 used the latest PCSP ideas, the core technique used, however, is this analytical approach.

A less extensive line of work proceeds via combinatorial gadgets that are analyzed using ideas based on the chromatic number of Kneser graphs and similar results. The first exemplar of this approach was the hardness of $(k-1)$ -coloring 2-colorable 3-uniform hypergraphs shown in [DRS05]. Unlike the analytic results for k -uniform hypergraphs mentioned above, this result does not show hardness of finding large independent sets. (This was later shown in [KS14] using the analytic approach, albeit under the $k=1$ conjecture.) A few recent results have revived this combinatorial approach, by re-deriving and improving some previous hardness results for hypergraph coloring using simpler proofs [Bha18; ABP19].

The third and most recent line of work adapts the universal algebraic method behind the complexity classification of constraint satisfaction problems that culminated in the resolution of the Feder-Vardi CSP dichotomy conjecture [Bul17; Zhu20]. Here, the coloring problem is viewed as a Promise Constraint Satisfaction Problem (PCSP), and its associated polymorphisms are then

analyzed.¹ In the cases when the polymorphisms are severely limited, one can show hardness via a reduction from Label Cover. The approach to study PCSP using polymorphisms originated in [AGH17], and was used to show hardness results for graph and hypergraph coloring in [BG16]. The algebraic theory was further developed significantly in [Bar+21] leading among other results to a proof of NP-hardness of coloring 3-colorable graphs.

In this work, we follow the algebraic approach to prove Theorem 35. Note that rainbow coloring is a natural Promise Constraint Satisfaction Problem. As proved in [Bar+21; BG21a], as with normal CSP, the complexity of a promise CSP is captured by its associated polymorphisms. Recall that polymorphisms of a PCSP are ways to combine multiple solutions of an instance satisfying the stronger predicate to obtain a solution to the instance satisfying the weaker predicate. The high level principle behind the algebraic approach is that the problem should be easy when it has a rich enough set of polymorphisms that include functions with strong symmetries, and hard when all its polymorphisms are somehow skewed and lack symmetries. This has been fully established for CSPs — when there are polymorphisms which obey weak near-unanimity, the CSP is polytime solvable, and otherwise NP-complete. The hardness part of this dichotomy is easier and was known for a while; the much harder algorithmic part was established only recently in [Bul17; Zhu20].

For promise CSPs, which form a much richer class, our current understanding is rather limited, for both the algorithmic and hardness sides. It is not clear (to even conjecture) what kind of lack in symmetries in the polymorphisms might dictate hardness, and how one might show the corresponding hardness. A simple (but rather limited) sufficient condition for hardness is when all the polymorphisms are dictators that depend on a single coordinate. In [AGH17], it has been proved that if all the polymorphisms of a PCSP are juntas, then the PCSP is NP-hard. This is the basis of the hardness results for $(2r+1)$ -SAT [AGH17] and 3-coloring graphs that admit a homomorphism to C_k for any fixed odd integer k [KO19]. The recent hardness of coloring 3-colorable graphs in [Bar+21] proceeds by showing that the absence of polymorphisms with the so-called C_k symmetry implies NP-hardness, and then verifying that 5-coloring lacks such polymorphisms.

It turns out that the polymorphisms of rainbow coloring can have low symmetries and be non-juntas. We will get around this by proving that these polymorphisms are C_k -free in the sense that there exists a constant number of coordinates and an assignment to them such that if we fix these coordinates to the assignment, the value of the function is fixed. This is also studied as certificate complexity in Boolean function analysis [APV16]. We then prove that if the polymorphisms of a PCSP are C_k -free, then the PCSP is NP-hard.

In order to prove that the polymorphisms have low certificate complexity, we use the connection between sensitivity and certificate complexity of functions. These two ways of characterizing

¹The proof in [DMR09] also implicitly studies polymorphisms and proves that they must have a small number of coordinates with sizeable influence and thus are not too symmetric. This influence-type characterization interfaces better with Unique Games or other highly structured forms of Label Cover.

²For the case of Boolean CSPs, the CSP is hard iff all polymorphisms are unary, i.e., either the dictator function or its complement.

³A C -junta is a function that depends on at most C inputs

complexity of functions are well studied in the context of Boolean functions. It is worth emphasizing that for our purposes, all we need is to show that low sensitivity (even sensitivity $2 \log C$) implies constant certificate complexity, and thus we are not interested in optimal gaps between sensitivity and certificate complexity. The famous sensitivity vs. block sensitivity conjecture [Nis89] states that these two parameters are in fact polynomially related. In one of the earliest works related to this problem, Simon [Sim83] proved that certificate complexity is at most exponential in sensitivity. We extend this to larger domains, and then use it to prove that the polymorphisms that we study have low certificate complexity.

The second step is to then use the fixing property to show NP-hardness of the PCSP. This is done by the usual paradigm of reducing from Label Cover using polymorphism tests (better known as long code tests) of functions associated with vertices of the Label Cover instance. A more structured form of the fixing property, where the C variables are fixed to the same value, is used in [BG21a] to show NP-hardness of certain Boolean PCSPs. However, in order to prove NP-hardness using our more general notion of fixing, we end up needing stronger properties of the Label Cover instance. As a result, our reduction is from the smooth Label Cover problem that was introduced and shown to be NP-hard in [Kho02a], and has found many applications in inapproximability since.

A natural question is to understand how far we can push these techniques. Our NP hardness reduction from smooth Label Cover works when the polymorphisms of the PCSP in hand are C -fixing for some constant C . As k increases, the polymorphisms of PCSP2-coloring a k -uniform hypergraph that is promised to be $(\epsilon - 1)$ -rainbow colorable get richer. When ϵ is at most $\frac{1}{6}$, the polymorphisms are C -fixing. At $k = 7$, we show that there is a polymorphism that is not C -fixing for any constant C . In fact, one would need C to be linear in the arity of polymorphisms which also rules out using smooth Label Cover with very strong soundness.

5.1.2 Prior work on rainbow coloring and related problems

Various notions of approximate coloring with rainbow colorability guarantees have been studied in the literature. Bansal and Khot [BK10] prove that when the input hypergraph is promised to be almost k -rainbow colorable, it is Unique Games hard to color it with $(1 - \epsilon)$ colors. Sachdeva and Saket [SS13] establish NP-hardness of coloring a k -uniform hypergraph when it is promised to be almost $\frac{k}{2}$ -rainbow colorable. This was extended by Guruswami and Lee [GL18] to perfectly $\frac{k}{2}$ -rainbow colorable hypergraphs. Guruswami and Saket [GS17] prove similar results assuming stronger forms of rainbow colorability in the completeness case. In [ABP20], Austrin, Bhangale and Potukuchi proved that it is NP-hard to color a k -uniform hypergraph when it is promised to be $(\epsilon - 2^{-k})$ rainbow colorable. On the other hand, when the hypergraph is promised to be $(\epsilon - \frac{1}{k})$ rainbow colorable, Bhattiprolu, Guruswami and Lee [BGL15] give algorithms to color the hypergraph with two colors that miscolors at most ϵ fraction of edges; this beats the 2^{k+1} fraction achieved by random coloring that is the best possible for general hypergraphs [Lis01]. Brakensiek and Guruswami [BG17] put forth a problem called label cover (to possibly serve as a perfect completeness variant surrogate for Unique games), and under

its conjectured inapproximability proved that it is hard to color a uniform $(k-1)$ -rainbow colorable hypergraph with $O(1)$ colors.

A related notion of hypergraph coloring is strong coloring where we color a uniform hypergraph with $s > k$ colors such that in any edge, all the vertices are colored with distinct colors. Brakensiek and Guruswami [BG16] prove that it is NP-hard to color a k -uniform hypergraph that is promised to be strongly colorable with s colors. Assuming the V Label Cover conjecture, it is hard to color $(1-\epsilon)$ -color k -uniform hypergraphs with strong chromatic number at most $k + \frac{1}{\epsilon}$ [BG17].

5.1.3 Outline

We start with a few notations and definitions in Section 5.2. In Section 5.3, we study polymorphisms of rainbow coloring. We first prove a result on sensitivity and certificate complexity and use it to prove properties of polymorphisms of the PCSP that we are studying. Then, we use these in Section 5.4 to prove NP hardness. Finally, we use the rainbow coloring of hypergraphs ideas to show hardness for Vector Bin Covering problem.

5.2 Preliminaries

5.2.1 Rainbow Coloring PCSP

In RAINBOW($k; r; q$) problem, the input is a uniform hypergraph. The goal is to distinguish between the cases when the hypergraph is rainbow colorable with r colors and when it cannot be rainbow colorable with q colors. More formally, we can define the problem as below:

Definition 37. (RAINBOW($k; r; q$)) In the RAINBOW($k; r; q$) promise CSP, $q \leq r \leq k$, we have a pair of predicates $(A; B)$ defined as follows:

- $A : [r]^k \rightarrow \{0, 1\}$; $1g : A(x_1; x_2; \dots; x_k) = 1$ if and only if $x_1; x_2; \dots; x_k \in [r]$.
- $B : [q]^k \rightarrow \{0, 1\}$; $1g : B(y_1; y_2; \dots; y_k) = 1$ if and only if $y_1; y_2; \dots; y_k \in [q]$.

Note that we need $q \leq r$ to be at most k since we cannot rainbow color a k -uniform hypergraph with more than k colors. We also need the condition that $r \leq k$ for the promise problem to make sense: If the hypergraph is rainbow colorable, we can infer that it is already r -rainbow colorable too. Thus, the promise problem is to identify if the hypergraph is rainbow colorable or it cannot even be rainbow colorable with q colors. Furthermore, in this work we will be only dealing with near perfect completeness case when hypergraph is $(1-\epsilon)$ -partite i.e. $r = k - 1$.

We now direct our attention to polymorphisms of RAINBOW($k; r; q$). By definition, the polymorphisms of hypergraph coloring PCSPs turn out to be colorings of certain tensor product hypergraphs. Fix n to be arity of the polymorphisms. We can infer that the polymorphisms of RAINBOW($k; r; q$) are proper q -rainbow colorings of the following k -uniform hypergraph $R_H(k; r)$:

- The vertex set of hypergraph is the set $V = [r]^n$.

- A k -element set $(v_1; v_2; \dots; v_k)$; where each $v_i \in [r]^n$ is an edge if and only if for every $j \in [n]$; the set $(v_1)_j; (v_2)_j; \dots; (v_k)_j$ is equal to $[r]$.

That is, a set of k vectors from $[r]^n$ forms an edge if in the matrix obtained by plugging these vectors as rows, all the elements from $[r]$ occur in every column.

5.2.2 Complexity measures of functions

Finally, we define the notions of sensitivity and C -ing of functions.

Definition 38. (Sensitivity) For a function $f : [r]^n \rightarrow [q]$, and an input $x \in [r]^n$, the sensitivity of f at x , denoted by $S(f; x)$ is defined as the number of coordinates such that changing at i can change the value of f . i.e. $S(f; x) = |\{i \in [n] : f(x) \neq f(x : x_i = a)\}|$.

Definition 39. (Sensitivity) The sensitivity of a function $f : [r]^n \rightarrow [q]$, denoted by $S(f)$ is defined as the maximum sensitivity over all x in $[r]^n$ i.e. $S(f) = \max_x S(f; x)$.

Definition 40. (C -ing) A function f from $[r]^n$ to $[q]$ is said to be C -ing for some integer C if there exists a set $S = \{s_1; s_2; \dots; s_C\} \subseteq [n]$ and a vector $\alpha = (\alpha_1; \alpha_2; \dots; \alpha_C) \in [r]^m$ such that $f(x) = p$ whenever $x_{s_i} = \alpha_i$ for all integers $1 \leq i \leq C$, for some $x \in [r]^n$.

5.3 Polymorphisms

In this section, we will analyze the properties of polymorphisms of rainbow coloring. In order to do so, we will prove that low sensitivity implies low certificate complexity. Using this, we will establish that the polymorphisms $\text{RAINBOW}(k; k-1; d^{\frac{k-2}{2}})$ are C -ing. Along the way, we will study rainbow colorings of various hypergraphs related to $\text{RH}(k; r)$. Finally, we will show that our techniques cannot prove hardness of $\text{RAINBOW}(7; 6; 2)$ by presenting a polymorphism that is not C -ing for any constant C .

5.3.1 Sensitivity vs certificate complexity

We extend a lemma of [Sim83] that proves that if a function has low sensitivity then the function is C -ing, to larger domains. The proof is along the same lines as the original proof.

Lemma 41. Let $f : [r]^n \rightarrow [q]$ be a function with sensitivity s , and let $b \in [q]$ such that $f^{-1}(b)$ is non-empty. Then $|f^{-1}(b)| \geq r^{n-s}$.

Proof. Fix s , and induct on n . The case $n = s$ is trivial. Let $x \in [r]^n$ be such that $f(x) = b$. Since $s < n$, there is a coordinate i that is not sensitive. Without loss of generality, let it be 1 and let $x = (x_1; y)$. As the first coordinate is not sensitive for f , we can conclude that $f(x; y) = b$ for all $x_1 \in [r]$.

Consider the set of functions $g : [r]^{n-1} \rightarrow [q]$; $g(u) = f(i; u)$; $i \in [r]$. Note that for each such g , the set $g^{-1}(b)$ is non-empty. In addition, for every $i \in [r]$, sensitivity of g_i is at most the sensitivity of f . Thus, by induction, we know that each g_i has at least r^{n-1-s} elements in

$[r]^{n-1}$ such that $f_i(u) = b$. Note that every such u gives $f(i; u) = b$. By combining over all s , we can conclude that there are at least $n-1 \cdot s = r^{n-s}$ elements $x \in [r]^n$ such that $f(x) = b$. \square

Lemma 42. Let $f : [r]^n \rightarrow [q]$ be a function with sensitivity $s < n/2$. Then, it is C -sensitive for $C = s(r-1)r^{2s+1}$.

Proof. We will actually prove a stronger statement that f is a C -junta. Let the set A denote the set of coordinates with non-zero influence, i.e. the coordinates that are sensitive for some input. Our goal is to upper bound the cardinality of A .

For a function $f : [r]^n \rightarrow [q]$, let the set of sensitive edges $E(f)$ be defined as the set of pairs of elements $x, y \in [r]^n$ such that $f(x) \neq f(y)$, and x, y differ on exactly one coordinate. From the sensitivity bound on f , we can deduce that

$$|E(f)| \leq s(r-1)r^n \tag{5.1}$$

Fix an arbitrary coordinate $i \in A$. There are elements $x, y \in [r]^n$ such that $x_i = a, y_i = b, a \neq b$; $f(x) \neq f(y)$, and x, y differ only in i th coordinate. Define a function $g : [r]^n \rightarrow \{0, 1\}$ as $g(z)$ is 1 if and only if $f(i; z) = f(x)$, and $f(i; z) = f(y)$ where we use the notation $f(i; z)$ to denote the vector f on $[r]^n$ obtained by inserting i in i th position to $z \in [r]^{n-1}$. Now, since $f(i; z)$ and $f(i; z)$ are both sensitive to at most s coordinates, $g(z)$ is sensitive to at most $2s$ coordinates. Also note that $g^{-1}(1)$ is non-empty. Thus, by Lemma 41, we can conclude that $g^{-1}(1)$ has at least r^{n-1-2s} elements. In other words, each sensitive coordinate contributes at least r^{n-1-2s} edges to $E(f)$. Thus, we can conclude that

$$|E(f)| \geq |A| r^{n-2s-1} \tag{5.2}$$

Combining Equation (5.1) and Equation (5.2), we get

$$|A| \leq s(r-1)r^{2s+1} \tag{5.3}$$

which proves the required claim. \square

5.3.2 Low sensitivity polymorphisms of rainbow coloring

We now turn our attention towards our main goal in this section: to show that polymorphisms of $\text{RAINBOW}(k; k-1; q)$ are C -sensitive for $C = d_{k-2}^{k-2} e$. As we have already mentioned earlier, the polymorphisms of rainbow coloring themselves are rainbow colorings of certain tensor product hypergraphs. To be precise, every polymorphism of $\text{RAINBOW}(k; r; q)$ are precisely q -rainbow colorings of $\text{RH}_k(k; r)$. Thus our new goal is to prove that any integer $k \geq 2$, any q -rainbow coloring of $\text{RH}_k(2q+2; 2q+1)$ is a C -sensitive function.

In order to achieve this, we will first define certain hypergraphs similar to $\text{RH}_k(k; r)$.

Definition 43. $H_h(r; s) = (V; E)$ is a r -uniform hypergraph where the vertex sets equal to $[r]^n$. A set of vectors $(u_1; u_2; \dots; u_r)$ is an edge if and only if

1. In every coordinate $i \in [n]$, at least $r-1$ elements occur i.e. $\sum_{j=1}^r (u_j)_i \geq r-1 \forall i \in [n]$.

2. All the r elements occur in at least s coordinates i.e. $\sum_j (u_j)_i = r$ for at least s choices of i in $[n]$.

The reason behind studying these hypergraphs is that q -rainbow colorings of $\mathcal{H}_h(2q+2; 2q+1)$ are very closely related to q -rainbow colorings of $\mathcal{H}_h(2q+1; c)$ for any absolute constant c . In fact if we can prove that q -rainbow colorings of $\mathcal{H}_h(2q+1; c)$ are C -sensitive, it implies that q -rainbow colorings of $\mathcal{R}_h(2q+2; 2q+1)$ are $\max(C; c)$ -sensitive. This is formally proved in Lemma 47. Thus our modified objective is to argue that q -rainbow colorings of $\mathcal{H}_h(2q+1; c)$ are C -sensitive. In order to do so, we inductively relate q -rainbow colorings of $\mathcal{H}_h(t; c)$ and $\mathcal{H}_h(t-1; c-1)$. As a base case, we have the following lemma:

Lemma 44. For all integers $q \geq 2$ and $n \geq 1$, the hypergraph $\mathcal{H}_h(2q-1; 1)$ cannot be rainbow colored with q colors.

Proof. We will use induction on q . For the case $q = 2$, rainbow coloring with 2 colors is the same as proper coloring the hypergraph with 2 colors. The fact that $\mathcal{H}_h(3; 1)$ cannot be two colored follows from [ABP20] (Lemma 3.2 with $d = 3$).

Suppose for contradiction that $\mathcal{H}_h(2q-1; 1)$ has a valid q -rainbow coloring. Let $r = 2q-1$ denote the uniformity of the hypergraphs. Consider the set of vectors in $[r]^n : f_i(i; i; \dots; i) = g$. As there are at most $q < r$ colors, some two elements of this set should have the same value. Without loss of generality, let $(r-1; r-1; \dots; r-1) = f$ and $(r; r; \dots; r) = g$ for some $c \in [q]$. Consider the $(q-2)$ -uniform hypergraph $\mathcal{H} = \mathcal{H}_h(r-2; 1)$. Note that every edge e in \mathcal{H} together with $u = (r-1; r-1; \dots; r-1)$ and $v = (r; r; \dots; r)$ forms an edge in $\mathcal{H}_h(r; 1)$. Thus, all the $q-1$ colors in $[q] \setminus \{c\}$ occur in every edge of coloring of $\mathcal{H}_h(r-2; 1)$ using f . This implies that we can get a valid $(q-1)$ -rainbow coloring of $\mathcal{H}_h(r-2 = 2(q-1)-1; 1)$ by restricting f to $[r-2]^n$, and replacing the color c using arbitrary color from $[q] \setminus \{c\}$. However, by induction hypothesis such a coloring cannot exist, and thus we have arrived at contradiction. \square

Now, we will use this to argue about q -rainbow colorings of $\mathcal{H}_h(2q+1; 3)$ via q -rainbow colorings of $\mathcal{H}_h(2q; 2)$. Consider the hypergraph $\mathcal{H}_h(2q; 2)$. A trivial way to q -rainbow color this hypergraph is to pick a coordinate $\ell \in [n]$, and partition the set $[2q]$ into q disjoint sets of size two, let's say $A_1; A_2; \dots; A_q$ and assign the value $p \in [q]$ to $f(x)$ if and only if $x_{\ell} \in A_p$. It turns out that this is the only way to q -rainbow color $\mathcal{H}_h(2q; 2)$. We prove it in the lemma below:

Lemma 45. Let f be a q -rainbow coloring of $\mathcal{H}_h(r = 2q; 2)$. Then, there exists an index $\ell \in [n]$, sets $A_1; A_2; \dots; A_q \subseteq [r]$ each of size 2, and mutually disjoint such that $f(x) = j$ iff $x_{\ell} \in A_j$.

Proof. First we will prove that the sensitivity of f is at most 1. Let $x = (x_1; x_2; \dots; x_n)$ be an arbitrary vector in $[r]^n$. Consider a $(q-1)$ -uniform hypergraph $\mathcal{H}(x)$ defined on $([r] \setminus \{x_1\}) \times ([r] \setminus \{x_2\}) \times \dots \times ([r] \setminus \{x_n\})$. We add a $(q-1)$ vector set as edge of $\mathcal{H}(x)$ if and only if it has at most one coordinate where there are missing elements i.e. all the x_i occur in all but one coordinate ℓ and in that coordinate, at most one value is missing.

Note that $\mathcal{H}(x)$ is isomorphic to $\mathcal{H}_h(2q-1; 1)$. From Lemma 44, we know that $\mathcal{H}(x)$ cannot be rainbow colored with q colors. Thus, when we view f as a coloring of $\mathcal{H}(x)$, there is an edge that has a color missing. Let it be denoted by $(y_1; y_2; \dots; y_{q-1})$. Let j be the coordinate

where there is a missing element. If there is no coordinate with missing element, can be arbitrary. Without loss of generality, let color $[q]$ be missing in e . Note that $xg[e]$ is an edge of $H_h(r; 1)$, and thus an edge of $H_h(r; 2)$ as well. Since f is a proper q -rainbow coloring of $H_h(2q; 2)$, we can conclude that $f(x) = 1$. In fact, we can actually deduce something stronger. Let $y \in [r]^n$ such that x and y differ on exactly one coordinate $j \in [n]$. Note that $yg[e]$ is also a valid edge of $H_h(2q; 2)$ since it has at most two coordinates where there are missing elements $i \in [q]$ and j . Thus, $f(y) = 1 = f(x)$. Thus, for every x , in except for one coordinate, changing the value of the coordinate preserves the value of f . In other words, the sensitivity of f is at most 1.

Using this, we will now prove that f is a dictator. Let i be an influential coordinate of f , i.e. there exists $x, y \in [r]^n$ differing only in i th coordinate such that $f(x) \neq f(y)$. We claim that $f(u) = f(x)$ for all $u \in [r]^n$ such that $u_i = x_i$, and $f(u) = f(y)$ if $u_i = y_i$. We will prove this by induction on the number of coordinates in which u and x differ excluding i . Since f has sensitivity at most 1, the only sensitive coordinate of f and x is i . Thus, for any u differing only in one coordinate from x (other than i) such that $u_i = x_i$ or y_i will have corresponding f value. Suppose that the statement holds for u differing from x in t coordinates excluding

Now, let u differ from x in $t + 1$ coordinates excluding i . We can find $v, w \in [r]^n$ such that v and x differ in t coordinates excluding i , $v_i = x_i$; w and y differ in t coordinates excluding i , $w_i = y_i$, and one of v and w differs from u in at most one coordinate. By induction hypothesis, $f(v) = f(x)$; $f(w) = f(y)$. Since v and w differ in a single coordinate i , i is the only sensitive coordinate of v and w . Thus, $f(u)$ is equal to either $f(v)$ or $f(w)$ depending on $u_i = x_i$ or y_i . This completes the inductive proof.

To complete the proof that f is a dictator, we will use this to show that there cannot be two influential coordinates. Suppose that there are two influential coordinates i_1, i_2 . From previous argument, we can infer that there are assignments $s_1, s_2 \in [r]$ such that assigning these to corresponding coordinates gives the value of f . Also note that assigning s_1 and s_2 to different values. Similarly, assigning s_2 and s_1 to different values. This gives rise to contradiction since if we set coordinate i_1 , f should be fixed irrespective of s_2 is equal to s_1 or s_2 . Thus, there can be only one influential coordinate for f or in other words f is a dictator.

Let p be the dictator coordinate of f , i.e. there exists a function $g: [r] \rightarrow [q]$ such that $f(x) = g(x_p)$. From the definition of the hypergraph $H_h(r; 2)$, for every $j \in [r]$, the set $S_j = \{g(x_i) \mid g(x_i) \neq g(x_j)\}$ should be equal to $[q]$. This proves that there exists sets $A_1, A_2, \dots, A_q \subseteq [r]$ each of size two, and mutually disjoint such that $g(x) = j$ if and only if $x \in A_j$, which proves the required claim. \square

We finish the chain of inductive arguments by proving a key property of rainbow colorings of $H_h(2q + 1; 3)$.

Lemma 46. Let $f: [2q + 1]^n \rightarrow [q]$ be a q -rainbow coloring of $H_h(r = 2q + 1; 3)$. Then, there exists an index $i \in [n]$, and $s \in [r]$ such that $S(f; x) = s$ for all $x \in [r]^n$ such that $x_i = s$.

Proof. Let $x = (x_1, x_2, \dots, x_n) \in [r]^n$ be an arbitrary vector in $[r]^n$. Similar to the previous lemma, we define the complement hypergraph associated with f . Consider a $(q - 1)$ -uniform

hypergraph $H(x)$ defined on $([r] \times x_1) \cup ([r] \times x_2) \cup \dots \cup ([r] \times x_n)$. We add a $(r-1)$ vector set as edge of $H(x)$ if and only if it has at most two coordinates where there are missing elements i.e. all the $[r] \times x_i$ occur in all but two coordinates and in these two coordinates, at least 2 values occur. Note that $H(x)$ is isomorphic to $H(r-1; 2)$.

We can view $f: [2q+1]^n \rightarrow [q]$ as a q -rainbow coloring of $H(x)$. If f is not a valid q -rainbow coloring of $H(x)$, by the same argument as in Lemma 45, we can deduce $S(f; x) \geq 2$. If f is a valid coloring of $H(x)$, we will use the properties proved in Lemma 45. Let us define a function $g: [r]^n \rightarrow [n] \cup \{?\}$ such that for a vector $x \in [r]^n$,

1. If f is a valid q -rainbow coloring of $H(x)$, then Lemma 45 implies that there exists a coordinate $i \in [n]$ such that f is a dictator in i th coordinate in $H(x)$. In this case, set $g(x) = i$.
2. If f is not a valid q -rainbow coloring of $H(x)$, let $g(x) = ?$.

First, we will prove that there exists an index $i \in [n]$ such that $g(x) = i; ?g$ for all $x \in [r]^n$. Suppose $g(x) = i \in [n]$, and $g(y) = j \in [n]$ where $x, y \in [r]^n$ and $i \neq j$. Since $g(x) = i$, there exist sets $S_1, S_2, \dots, S_n \subseteq [r]$ such that f is a dictator on i th coordinate in $S = S_1 \cup S_2 \cup \dots \cup S_n \subseteq [r]^n$. In particular, there is a subset $A \subseteq S_i$ such that $|A| = 2$, and $f(x), x \in S$, is equal to 1 if and only if $x_i \in A$. Similarly, there exist sets $T_1, T_2, \dots, T_n \subseteq [r]$ such that f is a dictator on j th coordinate in $T = T_1 \cup T_2 \cup \dots \cup T_n \subseteq [r]^n$. There exists a subset $B \subseteq T_j$ such that $|B| = 2$; and $f(x); x \in T$ is equal to $c \in [q]$ if and only if $x_j \in B$ for some $c \in [q]$. Combining the both, let $U_i = S_i \setminus T_i, U_j = T_j \setminus S_j, U_i \cap U_j = \emptyset, i \neq j \in [n]$. We can deduce that f is a dictator in both i and j coordinates in $U = U_1 \cup U_2 \cup \dots \cup U_n$. This implies that f is a constant function in U . Recall that there are two assignments s_i that make f equal to 1 and two assignments t_j that make f equal to $c \in [q]$. Thus, $f(x^0)$ is equal to 1 for some $x^0 \in U$ and $f(y^0) = c \in [q]$ for some $y^0 \in U$. This contradicts the fact that f is a constant function in U . Thus, there exists an index $i \in [n]$ such that $g(x)$ is either equal to i or is equal to $?$ for all $x \in [r]^n$. Without loss of generality let that be the first coordinate i.e. for all $x \in [r]^n, g(x) = i; ?g$.

Consider the case when $g(x) = ?$ for every $x \in [r]^n$. In this case, we know that $S(f; x) \geq 2$ for all $x \in [r]^n$. In particular, we can set arbitrary and say that $S(f; x) \geq 2$ whenever $x_1 = ?$. So we are only left with the case when there exists $\alpha \in [r]$ such that $g(x) = \alpha$. We will now prove that there exists $\beta \in [r]$ such that $g(x) = \beta$ whenever $x_1 = \alpha$, thus proving the required sensitivity bound.

Suppose for contradiction that for every $\beta \in [r]$, there exists $x \in [r]^n$ such that $x_1 = \alpha$, and $g(x) = \beta$. Consider a pair $u, v \in [r]^n$ such that $u_1 = \alpha; v_1 = \beta$ and $g(u) = g(v) = 1$. Let $u = (u_1; u_2; \dots; u_n), S_i = [r] \times u_i$ and f is dictator on i st coordinate in $S = S_1 \cup S_2 \cup \dots \cup S_n$. There is a function $h_1: S_1 \rightarrow [q]$ such that $f(x) = h_1(x_1)$ if $x \in S$ and $|h_1^{-1}(c)| = 2 \leq 8c \leq 2|q|$. Similarly, let $v = (v_1; v_2; \dots; v_n), T_i = [r] \times v_i$ and f is dictator on i st coordinate in $T = T_1 \cup T_2 \cup \dots \cup T_n$. There is a function $h_2: T_1 \rightarrow [q]$ such that $f(x) = h_2(x_1)$ if $x \in T$ and $|h_2^{-1}(c)| = 2 \leq 8c \leq 2|q|$. Let $U_i = S_i \setminus T_i$. Note that $U = U_1 \cup U_2 \cup \dots \cup U_n$ is non empty and f is dictator on i st coordinate in U as well. Note that $|U_i| \leq r-1$ for all $i \in [n]$. Thus, we can conclude that if $\beta \in U_1$, then $h_1(\beta) = h_2(\beta)$.

Applying this to all pairs u, v such that $g(u) = g(v) = 1$, we can infer that there exists a function $h : [r] \rightarrow [q]$ that satisfies the property that for all $x \in [r]^n$ such that $g(x) = 1$, let $x = (x_1; x_2; \dots; x_n)$, $S_i = [r] \times x_i$, $S = S_1 \cup S_2 \cup \dots \cup S_n$, then $f(y) = h(y_1)$ for all $y \in S$. As $r = 2q + 1 > 2q$, there exists $b \in [q]$ such that $|h^{-1}(b)| \geq 3$. Let $x \in [r]^n$ be such that $h(x_1) \in b$. From our assumption that for every $x \in [r]^n$ there exists $y \in [r]^n$ such that $g(y) = 1$ and $x_1 = y_1$, there exists $u \in [r]^n$ such that $u_1 = x_1$ and $g(u) = 1$. Now, let $u = (u_1; u_2; \dots; u_n)$, $S_i = [r] \times u_i$, $S = S_1 \cup S_2 \cup \dots \cup S_n$, and we know that $f(x) = h(x_1)$ if $x \in S$, and $h^{-1}(c) \setminus S_1 = 2q - 1$. However, this contradicts the fact that $h(u_1) = h(x_1) \in b$, and $|h^{-1}(b)| = 3$. Thus, there exists $x \in [r]^n$ such that $g(x) = 0$ for all $x \in [r]^n$ such that $x_1 = b$. \square

Finally, we will use the previous hypergraph coloring properties to argue about polymorphisms of rainbow coloring.

Lemma 47. Suppose $f : [2q+1]^n \rightarrow [q]$ is an n -ary polymorphism of $\text{RAINBOW}(2q+2; 2q+1; q)$ i.e. f is a proper q -rainbow coloring of $\text{RH}(2q+2; 2q+1)$. Then, there exist constant $C = C(q)$ independent of n such that f is C -xing.

Proof. Let $r = 2q + 1$. Let $f : [r]^n \rightarrow [q]$ be a polymorphism of $\text{RAINBOW}(2q + 2; 2q + 1; q)$. We can view f as a q -rainbow coloring of $\text{H}_h(r; 3)$ as the vertex set of $\text{RH}(r + 1; r)$ and of $\text{H}_h(r; 3)$ is equal to $[r]^n$. If it is not a valid q -rainbow coloring, there is an edge in which not all q colors appear. Let that edge be $e = (v_1; v_2; \dots; v_r) \in \text{RH}(r + 1; r)$ and $c \in [q]$ be a missing color in $f(v_1); f(v_2); \dots; f(v_r)$. Since this edge is part of $\text{H}_h(r; 3)$, except for 3 values of i , for all other i , the set $(v_1)_i; (v_2)_i; \dots; (v_r)_i$ is equal to $[r]$. Let the missing coordinates be the set $S = \{i_1; i_2; i_3\}$. Now consider an element u of $[r]^n$ such that it has missing values in S . From the definition of $\text{RH}(r + 1; r)$, we can deduce that the set u is an edge of $\text{RH}(r + 1; r)$. Since f is a valid q -rainbow coloring of $\text{RH}(r + 1; r)$, $f(u)$ is equal to c . Note that this should hold irrespective of what values u has, in coordinates outside S . This proves that f is C -xing with $C = 3$.

On the other hand if f is a valid q -rainbow coloring of $\text{H}_h(r; 3)$; using Lemma 46, we can deduce that there exists an index $i \in [n]$, and $b \in [q]$ such that $S(f; x) \geq 2$ whenever $x_i = b$. Now, we can consider a function $g : [r]^{n-1} \rightarrow [q]$ which on an input $y \in [r]^{n-1}$, is equal to $f(x); x = y; x_i = b$ i.e. we first insert b in i th position to y and then apply f . Note that g has sensitivity at most 1. From Lemma 42, we can conclude that g is C -xing for $C = 2(r-1)r^5$. In other words g is fixed by assigning values to a set of indices. This implies that f is also $C^0 = C + 1$ -xing since we can first assign b to x_i , then use C -xing property of g . \square

5.3.3 High sensitivity polymorphism of $\text{RAINBOW}(7; 6; 2)$

We show that there exists a function $f : [6]^n \rightarrow \{0, 1\}$ that is a polymorphism of $\text{RAINBOW}(7; 6; 2)$ that is not C -xing for any constant C . We start with a dictator but add just enough noise such that the function still remains being a polymorphism, but it is no longer xing. Let $wt(x)$ denote the number of $i \in [n]; i > 1$ such that $x_i = 1$. Let $S \subseteq [6]^n$ denote the set of

$x \in [6]^n$ such that $\text{wt}(x) > \frac{2n}{3}$. Let $h : [6]^n \rightarrow \{0, 1\}$ be noise function defined below. For a given $x \in [6]^n$, we define $f(x)$ as follows:

1. If $x \in S$
 - (a) If $x_1 = 3$, $f(x) = 0$
 - (b) Else, $f(x) = 1$
2. Else $f(x) = h(x)$:

A choice of noise function that works is inverting the original function $f(x)$ is defined as 1 if and only if $x_1 = 3$.

Proposition 48. The function $f : [6]^n \rightarrow \{0, 1\}$ defined above is a polymorphism of $\text{RAINBOW}(7; 6; 2)$ and it is not C -xing for any $C < \frac{n}{3}$.

Proof. Any polymorphism of $\text{RAINBOW}(7; 6; 2)$ is a proper 2-rainbow coloring of $\text{RH}(7; 6)$. Recall that rainbow coloring with two colors is the same as standard hypergraph coloring with two colors.

Polymorphism. In any set of 7 vectors E in $[6]^n$ such that all the 6 elements occur in all the coordinates, at most two vectors can be in S . This is because, in any set of three vectors in S , there exists a coordinate in which all three values are equal. Thus, there are vectors $x \in S$ with $x_1 = 3$ and vector $y \in S$ such that $y_1 = 3$ in E , which together ensures that E is not monochromatic.

C -xing. Suppose that there exists a set $T \subseteq [n]$ and $(t_1; t_2; \dots; t_m) \in [6]^m$ such that $f(x) = b$ for all x such that $x_i = t_i$ for all $1 \leq i \leq m$, for some $x \in [6]^n$. We will prove that $|T| \leq \frac{n}{3}$. Suppose for contradiction that $|T| > \frac{n}{3}$. First, if $1 \in T$, we can set all coordinates outside T to be equal to 6 , and in this case $f(x) = x_1$, which cannot be fixed if $1 \in T$. Thus $1 \notin T$. Next, if all the coordinates outside T are all equal to 1 , then $f(x)$ is equal to noise function, which is different from the case when the rest are equal to 1 . Thus, if f is indeed a C -xing function, for the C -xing assignment, the value of f should be independent of the assignment to the coordinates outside T . However, that is not the case as the value of f changes when we set all the coordinates outside T to be 1 or 6 . \square

5.4 NP-Hardness

In this section, we will use the properties of polymorphisms proved so far to argue about NP hardness of rainbow coloring PCSP. We will prove the below theorem:

Theorem 49. Suppose that there exists a constant c such that for all integers $n \geq 1$, every n -ary polymorphism of $\text{RAINBOW}(k; k-1; q)$ is C -xing. Then, the corresponding decision problem $\text{RAINBOW}(k; k-1; q)$ is NP hard.

Before delving into the proof of Theorem 49, we first mention that this theorem together with Lemma 47 implies Theorem 35. In Lemma 47, we have proved that for every q , the polymorphisms of $\text{RAINBOW}(2q+2; 2q+1; q)$ are C -xing. This fact combined with Theorem 49

implies that $\text{RAINBOW}(2q+2; 2q+1; q)$ is NP hard for every $q \geq 2$. This already proves Theorem 35 when k is even. In order to prove when k is odd, note that we can use Lemma 45 in Lemma 47 to prove that the polymorphisms $\text{RAINBOW}(2q+1; 2q; q)$ are C- xing . We can combine this with Theorem 49 to prove Theorem 35 when k is odd.

The rest of this section is dedicated to proving Theorem 49. Like various other hardness of approximation results, we will use the standard label cover with long code framework. We reduce smooth label cover introduced in [Kho02a] to rainbow coloring PCSP.

First we formally define the gap Label Cover problem below:

Definition 50. ($(1; \gamma_{LC})$ Gap Label Cover) In $(1; \gamma_{LC})$ Gap Label Cover, we are given a Label Cover instance $(G = (L; R; E); \{c_e\}_e)$, and the goal is to distinguish between the following two cases:

1. There is a labelling $\sigma : G \rightarrow \Sigma$ that satisfies all the constraints.
2. No labelling can satisfy γ_{LC} fraction of constraints.

As mentioned earlier, we need stronger properties of the Label Cover instance that we are starting with. One such property is smoothness.

Definition 51. (Smoothness) A Label Cover instance $(G = (L; R; E); \{c_e\}_e)$ is said to be $(\gamma; J)$ smooth if for any vertex $s \in L$ and a set of labels $S \subseteq \Sigma$; $|S| = J$, over a uniformly random neighbor $v \in R$, $\Pr(\sum_{e=(s,v)} c_e(s) < \gamma |S|) \leq \epsilon$.

The following is an easier version of Theorem 1.17 in [Wen13].

Theorem 52. For every $\gamma; \gamma_{LC} > 0$ and $J \in \mathbb{Z}_+$, there exists $\epsilon = \epsilon(\gamma; \gamma_{LC}; J)$ such that $(1; \gamma_{LC})$ Gap Label Cover with Label size J that is promised to be $(\gamma; J)$ -smooth is NP hard.

Reduction. We start with $(1; \gamma_{LC})$ Gap Label Cover instance $(G = (L; R; E); \{c_e\}_e)$ that is promised to be $(\gamma; J)$ -smooth, for γ and γ_{LC} to be set later, and output a PCSP instance. Let $n = |L|$ denote the label size $n = |J|$. For each vertex $x \in L \cup R$, we add a set of nodes K_x of size $[kn - 1]^n$, indexed by n -length vectors. We add two types of constraints:

1. Coloring constraints: Inside every vertex of the Label Cover instance, we add the following constraints among the $[kn - 1]^n$ nodes. We add the constraint that the promise relation should be satisfied in the set of k nodes $\{u_1; u_2; \dots; u_k\} \in [kn - 1]^n$, if for every $i \in [n]$, the set $\{u_j(i); j \in [k]\}$ has cardinality $k - 1$.
2. Equality constraints: For every constraint $e : u \rightarrow v$ of the Label Cover, we add a set of equality constraints between nodes $K_u, v \in K_v$ if for all $i \in [n]$, $u_i = v_{e(i)}$.

Note that the Coloring constraints give rise to rainbow coloring uniform hypergraphs. It is yet unclear how we can justify adding equality constraints. One way to handle the equality constraints is to have a single node for all the vertices corresponding to equality constraint. However, this fails if we want to add a coloring constraint that involves two copies of the same vertex. A neater way to get around this is to argue that adding equality constraints does not change the set of polymorphisms, and thus the hardness of the predicate remains same with or without equality constraints. This simple fact is proved in Section 5.6.

Completeness If the label cover instance is satisfiable, then PCSP instance that is being output can be satisfied by assignment from Σ . Suppose $\sigma : L \rightarrow \Sigma$ is a labelling that satisfies

all the constraints of the Label Cover. For every vertex $u \in L \cup R$, we can assign the value $\phi(u)$. In other words, in every long code, we are assigning corresponding dictator function. The coloring constraints are defined precisely such that this assignment satisfies the constraints. The equality constraints also follows since the labelling satisfies all the constraints of the Label Cover.

Soundness If the Label Cover is not ϵ -satisfiable, we need to show that there is no assignment of the PCSP instance that satisfies all the constraints. Taking contrapositive, if there is an assignment ϕ to PCSP instance that satisfies all the constraints, then we will prove that there is an assignment to the Label Cover instance that can satisfy a fraction of constraints, for an absolute constant ϵ . Taking $\epsilon < c$, we can arrive at a contradiction, thus proving that there is no assignment ϕ to PCSP that satisfies all the constraints.

Let $\phi : [k-1]^n \rightarrow [q]$ denote the assignment to the PCSP instance that satisfies all the constraints for $v \in L \cup R$. From the coloring constraints, we can infer that ϕ is a n -ary polymorphism of $\text{RAINBOW}(k; k-1; q)$. Thus, it is ϵ -satisfying for a constant ϵ independent of n .

For every vertex $v \in L \cup R$ of the Label Cover instance, we will assign a set of labels $A(v)$. For vertices v in L , $A(v)$ is the ϵ -satisfying set. Since the Label Cover instance is smooth, we will only consider the constraints where all these labels go to distinct labels on the right under projections. We can set the smoothness parameter to be 0.1 for example, and we will be left with $\frac{9}{10}$ fraction of original constraints. We will prove that there is an assignment that satisfies a fraction of these constraints, for an absolute constant ϵ , which will prove the original soundness claim. Thus in all the remaining constraints, the set of labels $A(v)$ go to distinct labels on the right. Thus, for a vertex $v \in R$, each constraint $(u; v)$ gives rise to ϵ coordinates $\{u; v\}(A(u))$. For these vertices, we set $A(v)$ to be the set of maximal disjoint sets of such a projection ϕ of coordinates.

In order to prove that there is a good labelling to the Label Cover, we assign a label to every vertex v from $A(v)$ uniformly at random and prove that it satisfies constant fraction of constraints with non-zero probability. We will in fact show that the random assignment satisfies a constant fraction of constraints in expectation. We prove this in two steps. First, we show that for every constraint $(u; v)$ of the Label Cover, there exists $x \in A(u); y \in A(v)$ such that $\phi_{u;v}(x) = y$. This follows from the definitions of $A(v)$: suppose the projection $\phi(u)$ is disjoint from $A(v)$. In that case, we can add the projection $\phi(u)$ to $A(v)$ to get a larger set $A(v)$, which contradicts the fact that $A(v)$ is the maximal such union of disjoint projections. This implies that the uniformly random labelling satisfies each constraint $(u; v)$ of Label Cover with probability at least $\frac{1}{|A(u)||A(v)|}$.

To complete the proof, we need to bound the size of $A(v)$. As we have already mentioned, for $v \in L; |A(v)| \leq C$. We bound the size of $A(v)$ for vertices v in R using the below lemma.
Lemma 53. Suppose $\phi : [k-1]^n \rightarrow [q]$ is a polymorphism of $\text{RAINBOW}(k; k-1; q)$. Let $A_1; A_2; \dots; A_t$ be mutually disjoint subsets of $[k-1]^n$ such that each of them is ϵ -satisfying set of ϕ . Then, $t < k/\epsilon$.

Proof. First note that all the A_i 's should intersect to the same value b since otherwise, the vector $u \in [k-1]^n$ that has all the intersecting sets A_i 's is forced to be equal to multiple colors b at the same time. Let all the A_i 's be C -intersecting with respect to value $b \in [q]$ i.e. for each $i \in [n]$, there exists an assignment f_i such that if the value of x in A_i is equal to the assignment, then the value of $f(x)$ is equal to b irrespective of values of coordinates outside A_i . If $t \leq k$, we can find $u_1; u_2; \dots; u_k$ such that all $[k-1]$ occur in every coordinate, and u has the intersecting assignment of A_i . This implies that $f(u_i) = b$ for all i . However, note that $u_1; u_2; \dots; u_k$ is an edge of $R_H(k; k-1)$, and thus iff f is a polymorphism of $\text{RAINBOW}(k; k-1; q)$, all the $[q]$ elements should occur iff $f(u_1); f(u_2); \dots; f(u_k) = b$. This is a contradiction since for all $f(u_i) = b$. \square

From the lemma, we can infer that the cardinality $|A(v)|$ for $v \in R$ is at most kC . Combining this with the above, we can deduce that there is an assignment that satisfies a constant fraction of constraints, which is a constant fraction of constraints, independent of n .

5.5 Application: Vector Bin Covering

5.5.1 Problem overview

In the Bin Covering problem, the input is a set of items with sizes $a_1; a_2; \dots; a_n$. The objective is to partition them into a maximum number of parts such that in each part, the sum of the items is at least 1. The problem is a classic NP-Hard problem and admits a Polynomial Time Approximation Scheme (PTAS) [CJK01]. Vector Bin Covering is a multidimensional variant of Bin Covering. In this problem, the input is a set of 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.

Definition 54. (Vector Bin Covering) In the Vector Bin Covering problem, we are given 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

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

The problem is also referred to as “dual Vector Packing” in the literature. It is introduced by Alon et al. [Alo+98] who gave a $O(\log d)$ factor approximation algorithm. They also gave a 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, Ray [Ray21] showed that the 2-dimensional Vector Bin Covering problem is hard to approximate within a factor of $\frac{2}{3}$.

We show $\frac{\log d}{\log \log d}$ hardness, almost matching the $O(\log d)$ factor algorithm [Alo+98].

Theorem 55. d -dimensional Vector Bin Covering is NP-hard to approximate within $\frac{\log d}{\log \log d}$ factor.

5.5.2 Hardness of Vector Bin Covering via Rainbow Coloring

Our hard instances for Vector Bin Covering are when the vectors are from \mathbb{F}_2^d . In this setting, the Vector Bin Covering problem is closely related to the hypergraph rainbow coloring problem.

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

Lemma 56. Given a hypergraph $H = (V; E)$ and a parameter k , there is a polynomial time reduction that outputs a Vector Bin Covering instance $(V; \{v_1, \dots, v_n\}; \{0, 1\}^d)$ with $n = |V|$; $d = |E|$ such that

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

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

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

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

Proof. Let $n = |V|$; $d = |E|$. We order the edges as $E = \{e_1; e_2; \dots; e_d\}$. We output a set of vectors $V = \{v_1; v_2; \dots; v_n\}$ 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:

- (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,

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

- (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

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

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 $e = \{i_1, \dots, i_d\}$ of the hypergraph H . As $\prod_{j \in A_i} (v_j) \leq 1$ for all $i \in [2]$, there exist $v_1, v_2 \in e$ such that $v_1 \in A_1, v_2 \in A_2$. Thus, the coloring ϕ is a proper 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 55. Note that the dimension of the resulting vectors in the Vector Bin Covering instance is equal to the number of edges $|E|$ of the hypergraph H , and the gap in the optimal Bin Covering value is equal to $\frac{1}{k}$, the number of colors. Hence, to obtain better inapproximability results for Vector Bin Covering that grow with our goal is to show the hardness of approximate rainbow coloring on hypergraphs with edges where the number of colors is as large a function of n 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 $k = \frac{\log m}{\log \log m}$ colors.

Theorem 57. 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 $k = \frac{\log m}{\log \log m}$.

We defer the proof of Theorem 57 to Section 5.5.3.

We now prove the hardness of Vector Bin Covering using Theorem 57.

Proof of Theorem 55 Using Theorem 57 combined with the reduction in Lemma 56, we get that the following problem is NP-hard. Given a set of vectors $v_1, v_2, \dots, v_n \in \{0, 1\}^d$, distinguish between

1. V can be partitioned into $k = \frac{\log d}{\log \log d}$ parts such that in each part, the sum of vectors is at least 1 in every coordinate.
2. 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 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 with $k = \frac{\log d}{\log \log d}$. \square

5.5.3 Proof of Theorem 57

Our proof follows the same lines as that of Theorem 49. We present the full proof here for the sake of completeness.

We first need a slightly different notion of C - coloring for $C = 1$.

Definition 58. (1-coloring [BG16; GS20b]) A function $f : [k]^n \rightarrow \{0, 1\}$ is said to be 1-coloring if there exists an index $i \in [n]$ and values $a, b \in [k]$ such that

$$f(x) = 0 \iff x_i = a \quad \text{and} \quad f(x) = 1 \iff x_i = b$$

In the analysis of our reduction later, we need a definition and a lemma from [ABP20].
 Definition 59. (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 r -dimensional vectors over $[k]$ i.e. $V = [k]^n$. A set of k vectors $v^1; v^2; \dots; v^k$ form an edge of the hypergraph if

$$\sum_{i=1}^k v_i^j = r \quad \forall j \in [k]$$

Lemma 60. For every $k \geq 2$, the hypergraph $H_{b/2, c}^n[k]$ is not 2-colorable.

We analyze the polymorphisms of the underlying PCSP used in our reduction.

Lemma 61. Fix $k \geq 3$. Suppose $f : [k]^n \rightarrow \{0, 1\}$ satisfies the below two-coloring property: For every $2k$ vectors $v^1; v^2; \dots; v^{2k} \in [k]^n$ with

$$f(v_i^j) = 0 \iff i \in [k]$$

we have

$$f(v_i^j) = 1 \iff i \in [k]$$

Then, f is 1-coloring.

Proof. We first prove that there exist $a \in [n]; b \in [k], c \in \{0, 1\}$ such that $f(x) = c$ for all $x \in [k]^n$ with $x_i = a$. Suppose for contradiction that this is not the case. Then, for every $i \in [n]; j \in [k]$ there exist vectors $x^{ij}; y^{ij} \in [k]^n$ such that $x_i^{ij} = y_i^{ij} = j$, and $f(x^{ij}) = 0$ where $f(y^{ij}) = 1$.

Let $r = b/2, c$. 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 60), 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 $v^1; v^2; \dots; v^k \in [k]^n$ and $b \in \{0, 1\}$ such that $f(v^j) = b$ for all $j \in [k]$. Furthermore, there are at most missing values in these vectors i.e.

$$\sum_{i=1}^k v_i^j = r \quad \forall j \in [k]$$

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

1. $f(u^j) = b$ for all $j \in [r]$.
2. For every $j \in [n]$,

$$\sum_{i=1}^k v_i^j + \sum_{i=1}^r u_i^j = r \quad \forall j \in [n]$$

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

$$f(w^j) : j \in [2k] \text{ is } [k] \text{ is } [n]$$

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

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

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

$$f(v^j) : j \in [2k] \text{ is } [k]$$

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

We are now ready to prove Theorem 57. Our hardness result is obtained using a reduction from the Label Cover problem, similar to Theorem 49.

Reduction. We start with the Label Cover instance $C = (V = L \cup R; E = \{L = R\})$ from Theorem 128 and output a hypergraph $\mathcal{H} = (V^0, E^0)$. Let n denote the label size $n = |L| = |R|$. For each vertex $x \in L \cup R$, we have a long code containing a set of nodes of size $|k|^n$, indexed by n length vectors.

1. The vertex set of the hypergraph \mathcal{H} is the union of all the long code nodes.

$$V^0 = \bigcup_{v \in V} K_v$$

2. Edges of the hypergraph: For every vertex $v \in V$ of the Label Cover instance, we add an edge in E^0 for each set of $2k$ vectors $\{v^1, v^2, \dots, v^{2k}\}$ in K_v , if

$$f(v^j) : j \in [2k] \text{ is } [k] \text{ is } [n] \tag{5.4}$$

The number of edges in \mathcal{H} is at most

$$|E^0| \leq \sum_{v \in V} \sum_{\substack{K \\ |K|=2k}} |K|^{O(k)}$$

3. Equality constraints: For every constraint $u = v$ of the Label Cover, we add a set of equality constraints between nodes $x \in K_u, y \in K_v$ if for all $i \in [n]$, $x_i = y_{e(i)}$. 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. 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. $f(x) = f(y)$ for all pairs of nodes x, 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 $V \rightarrow [k]$ that satisfies all the constraints. We define the coloring $f : V \rightarrow [k]$ of H as follows. For every node $v \in K_v$, we set the dictatorship function

$$f(x) = x_{(v)}$$

By the constraints added in Equation (5.4), the function f is a valid k -rainbow coloring of H . As f satisfies all the constraints of the Label Cover, the coloring f satisfies all the equality constraints.

Soundness. Suppose that there is no labeling $V \rightarrow [k]$ that satisfies all the constraints in G . Then we claim that there is no coloring of H that respects all the equality constraints. Suppose for contradiction that there is a coloring $f : V \rightarrow [k]$ that respects all the equality constraints.

Consider a vertex $v \in V$. The function $f_v : [k]^n \rightarrow \{0, 1\}$, defined as $f_v(x) = 1$ if $x_{(v)} = f(v)$ and 0 otherwise, satisfies the conditions in Lemma 6.1. Thus, f_v is 1-satisfying for every $v \in V$. Hence, there is a function $L : V \rightarrow [k]$ such that for every $v \in V$, f_v is 1-satisfying on the coordinate $L(v)$. We now claim that the labeling $L : V \rightarrow [k]$ defined as $L(v) = L(v)$ satisfies all the constraints in G .

Consider an edge $e = (u, v)$, $u \in L$; $v \in R$ with the projection constraint $\pi_e : [k] \rightarrow [k]$. 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-satisfying property of f_u , we have $u_i \in [k]$ such that

$$f_u(x) = 1 \iff x_{L(u)} = u_i \quad \text{and} \quad f_v(x) = 1 \iff x_{L(v)} = v_j$$

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

$$f_v(y) = 1 \iff y_{L(v)} = v_j \quad \text{and} \quad f_u(y) = 1 \iff y_{L(u)} = u_i$$

By the equality constraints, $f_u(x) = f_v(y)$ for all $x, y \in [k]^n$ such that $x_i = y_{e(i)}$ for all $i \in [n]$. Let $y^0 \in [k]^n$ be an arbitrary vector with $y_{L(u)}^0 = u_i$; $y_{L(v)}^0 = v_j$. We choose $x^0 \in [k]^n$ such that for all $i \in [n]$, $x_i^0 = y_{e(i)}^0$. Note that $x_{L(u)}^0 = u_i$. Thus, $f_u(x^0) = 1$ while $f_v(y^0) = 0$. However, this contradicts the equality constraints.

5.6 Adding equality constraints

We will prove that adding equality relation structure does not affect the polymorphisms. By equality relation structure, we mean $(=; =) := (A \subseteq [q_1]^2; B \subseteq [q_2]^2)$ where $q_1 = q_2$ and $A = f(x; x) : x \in [q_1]^n; B = f(y; y) : y \in [q_2]^n$.

Lemma 62. Suppose $P = (A_1; B_1); (A_2; B_2); \dots; (A_m; B_m)$ is a PCSP template. Let the template Q be obtained by adding relational structure $(A^0; B^0) := (=; =)$ to P . Then, under log space reductions P is equivalent to Q .

Proof. We will show that P and Q have identical set of polymorphisms. Note that Q contains all the relations structures in P , polymorphisms of Q are a subset of P . We claim that the reverse direction also holds because every function is a polymorphism. Consider a polymorphism $f : [q_1]^n \rightarrow [q_2]$ be n -ary polymorphism of P . Consider vectors $v_1; v_2; \dots; v_n$ such that for all $i \in [n]; (v_i)_1; (v_i)_2 \in A^0$. Note that this implies that for all $i, (v_i)_1 = (v_i)_2$. Consider the tuple $((v_1)_1; (v_2)_1; \dots; (v_n)_1); f((v_1)_2; (v_2)_2; \dots; (v_n)_2) = f((v_1)_1; (v_2)_1; \dots; (v_n)_1); f((v_1)_1; (v_2)_1; \dots; (v_n)_1) \in B^0$. Thus, f is a polymorphism of $(=; =)$ as well which implies that f is a polymorphism of Q . It has already been shown [Pip02b; BG21b; Bar+21] that if polymorphisms of a PCSP are a subset of polymorphisms of Q then Q is log space reducible to P . Thus, P and Q are equivalent under log space reductions. \square

Chapter 6

Robust Algorithms and SDPs for Promise CSPs

6.1 Introduction

Horn-SAT and 2-SAT are Boolean constraint satisfaction problems (CSPs) that admit simple combinatorial algorithms for satisfiability. They are both examples of bounded width CSPs, which means that the existence of locally consistent assignments (which satisfy all local constraints involving some bounded number of variables, and which are consistent on the intersections) implies the existence of a global satisfying assignment.

While the simple local propagation algorithms for Horn-SAT and 2-SAT work when the instance is perfectly satisfiable, they are not robust to errors—if the given instance is almost satisfiable, i.e., $(1 - \epsilon)$ -satisfiable for $\epsilon > 0$, the local consistency based algorithms do not guarantee solutions that satisfy almost all the constraints. In a beautiful work, Zwick [Zwi98] initiated the study of finding “robust” algorithms for Constraint Satisfaction Problems (CSPs), namely algorithms that output solutions satisfying a $(1 - \epsilon)$ fraction of the constraints when the instance is guaranteed to be $(1 - \epsilon)$ -satisfiable, where $\epsilon > 0$ as $\epsilon \rightarrow 0$. Zwick obtained robust algorithms for 2-SAT using SDP rounding and for Horn-SAT based on LP rounding. The PCP theorem together with Schaefer's reductions [Sch78] shows that Boolean CSPs that are NP-Hard are also APX-hard with perfect completeness, which in particular means that they do not admit robust satisfiability algorithms. The only other interesting Boolean CSP besides Horn-SAT and 2-SAT for which satisfiability is polynomial-time decidable is Linear Equations modulo 2. Hastad [Has01] in his seminal work showed that even 3-LIN (when all equations involve just three variables), for every $\epsilon > 0$, it is NP-Hard to output a solution satisfying a $(1 - \epsilon)$ fraction of the constraints even when the instance is guaranteed to have a solution satisfying a fraction of the constraints.

Unlike Horn-SAT or 2-SAT, the satisfiability algorithm for 3-LIN is not local, and 3-LIN does not have bounded width. Thus, for Boolean CSPs, bounded width characterizes robust

¹For CSPs, this is equivalent to solvability by $O(1)$ rounds of the Sherali Adams LP hierarchy.

satisfiability. For CSPs over general domains, a landmark result in the algebraic approach to CSP due to Barto and Kozik [BK14b] showed that CSPs that are not bounded width can express linear equations. A reduction from Adami's result then shows that CSPs that are not bounded width do not admit robust algorithms. Guruswami and Zhou [GZ12] conjectured the converse—namely that all bounded width CSPs, over any domain, admit robust algorithms. Another work by Barto and Kozik [BK16] resolved this conjecture in the affirmative, thus giving a full characterization of CSPs that have robust algorithms.

In this work, we study robust algorithms for PCSPs. Broadly speaking, the study of PCSPs has been on two fronts: First, understanding which PCSPs can be solved in polynomial time, motivated by questions such as approximate graph coloring (and $\text{MAJ-}k\text{-SAT}$). Second, understanding the power of various algorithms for PCSPs. We initiate the study of robust algorithms for PCSPs motivated by both these directions. On one hand, robust algorithms are important on their own, understanding whether there are algorithms that work even with a small noise in the input. On the other hand, robust algorithms are equivalent to bounded width levels of Sherali Adams, and solvability by basic SDP for CSPs. The question of whether the same holds for PCSPs as well is a way to understand the power of these algorithmic tools themselves.

As is the case with CSPs, a natural approach to characterize which PCSPs have robust algorithms is via bounded width of PCSPs. However, it turns out that bounded width for PCSPs is weaker than having robust algorithms. Concretely, Atserias and Dalmau [AD22] have proved recently that the PCSP (1-in-3-SAT, NAE-3-SAT) does not have bounded width. As we shall show later, this PCSP indeed has a robust algorithm. Atserias and Dalmau also proved that this PCSP indeed can be solved by levels of Sherali-Adams, and as we shall later, it can also be solved using the basic SDP.

Having ruled out the connection to bounded width, we study robust algorithms for PCSPs directly via their polymorphisms. Polymorphisms are closure properties of satisfying solutions to (Promise) CSPs. As a concrete example, consider a 2-SAT CSP: given an instance of 2-SAT over n variables x_1, x_2, \dots, x_n , suppose that u, v, w are three assignments to these variables satisfying all the constraints. In then the assignment z that is coordinatewise Majority operation on three bits, i.e. $z_i = \text{MAJ}(u_i, v_i, w_i)$ for every $i \in [n]$, also satisfies all the constraints. This shows that the Majority function on three variables, or more generally, any odd number of variables is a polymorphism of the 2-SAT CSP. Similarly, the Parity function on any odd number of variables is a polymorphism of LIN, whereas there are no non-trivial polymorphisms for 3-SAT. Polymorphisms are the central objects in the universal algebraic approach to CSPs [JCG97; Jea98; BJK05], which has then been extended to PCSPs [BG21b; Bar+21].

At a high level, the existence of non-trivial polymorphisms implies algorithms, and vice-versa. The key challenge is to precisely characterize which polymorphisms lead to algorithms. It is known that the polymorphism family of a PCSP fully captures its computational complexity i.e., if there are PCSPs \mathcal{C}^0 such that the polymorphism family of \mathcal{C}^0 is contained in $\text{Pol}(\mathcal{C}^1)$, then \mathcal{C}^0 is formally easier than \mathcal{C}^1 , i.e., there is a gadget reduction from \mathcal{C}^0 to \mathcal{C}^1 . It turns out that this gadget reduction also preserves the existence of robust algorithms. This leads to the following questions: Which polymorphisms lead to robust algorithms for PCSPs? Can we use the

polymorphism characterization of robust algorithms to relate Basic SDP and robust algorithms for PCSPs?

We make progress on these questions on two fronts: first, for Boolean symmetric PCSPs where we allow negation of variables, we study which polymorphisms lead to robust algorithms, and which Boolean symmetric PCSPs do not admit robust algorithms. Second, towards understanding the power of basic SDP for promise CSPs, we introduce a minion and show that a PCSP can be solved by basic SDP if and only if there is a minion homomorphism from the minion of polymorphisms of.

As is the case with CSPs, if a PCSP is NP-Hard, then it does not admit robust algorithms. Thus, the above question is relevant only for PCSPs that can be solved in polynomial time. A large class of PCSPs for which polynomial time solvability has been fully characterized is the Boolean symmetric PCSPs. In [BG21b], the authors showed that for a Boolean symmetric PCSP folding i.e., we allow negating the variables (later this restriction was removed by [Fic+19]), can be solved in polynomial time if and only if it contains at least one of Alternate-Threshold (AT) Majority (MAJ) or Parity polymorphisms of all odd arities. While AT and MAJ are noise stable functions, Parity is highly sensitive to noise i.e., if we perturb each input with a small probability, the function output changes significantly. In fact, Parity having low noise sensitivity can also be viewed as one reason why LIN, despite having Parity as polymorphisms, does not admit robust algorithms. Thus, for Boolean symmetric folded PCSPs, a natural candidate characterization of robust algorithms is the existence of AT or MAJ polymorphisms.

On the algorithmic front, we prove that this indeed is the case, and on the hardness side, assuming the Unique Games Conjecture [Kho02a], we show that the absence of MAJ polymorphisms implies the lack of robust algorithms.

6.1.1 Robust algorithms

Our main algorithmic result is the following.

Theorem 63. Every Boolean folded PCSP that contains AT or MAJ polymorphisms of all odd arities admits a polynomial time robust algorithm. In particular,

1. If Γ contains MAJ polymorphisms of all odd arities, then for every $\epsilon > 0$, there exists a polynomial time algorithm that given an instance ϕ that is promised to have a solution satisfying $1 - \epsilon$ fraction of the constraints, outputs a solution satisfying $O(\frac{1}{3})$ fraction of the constraints³.
2. If Γ contains AT polymorphisms of all odd arities, then for every $\epsilon > 0$, there exists a polynomial time algorithm that outputs a solution satisfying $O(\frac{\log \log \frac{1}{\epsilon}}{\log \frac{1}{\epsilon}})$ fraction of the constraints on an instance promised to have a solution satisfying $1 - \epsilon$ fraction of the constraints.

²A predicate P is symmetric if for every satisfying assignment $(x_1; \dots; x_n)$ to P , any permutation of that assignment also satisfies P . For a Boolean predicate whether an assignment satisfies a predicate depends only on the Hamming weight. A PCSP is said to be symmetric if all the predicates in the template are symmetric.

³Here, O hides multiplicative poly logarithmic factors.

We obtain our robust algorithms by rounding the basic Semi Definite Programming (SDP) relaxation. For the Majority polymorphisms case, we use the same robust algorithm of Charikar, Makarychev, Makarychev [CMM09] for 2-SAT, with a completely different analysis. The exact version of 2-SAT has a simple algorithm based on rounding basic SDP: suppose that we have the predicate $x_1 \oplus x_2$. We find vectors v_0 and v_1, v_2 that satisfy the basic SDP constraints. As the basic SDP has zero error, we get that $\langle v_1, v_0 \rangle + \langle v_2, v_0 \rangle = 0$. This gives a simple rounding algorithm: we round x_i to True if and only if $\langle v_i, v_0 \rangle \geq 0$. This is not a robust algorithm: when there is a weaker guarantee that $\langle v_1, v_0 \rangle + \langle v_2, v_0 \rangle = \epsilon$, the above rounding can round both variables to False. Zwick [Zwi98] gave the first robust algorithm for 2-SAT where he does “outward rotation” where he rotates the vectors with $\langle v_i, v_0 \rangle \geq \frac{1}{3}$ by certain angle depending on $\langle v_i, v_0 \rangle$ before using the above algorithm. [CMM09] gave an algorithm that does the update continuously instead of a sudden jump at a fixed threshold, that gave an optimal error $\epsilon^{1/3}$. We use the same algorithm as theirs, but the analysis is completely different as we need our analysis to be predicate independent, and just use the existence of MAJ polymorphisms as a black box.

As a concrete example, consider the PCSP (2-in-4-SAT, 4-SAT). Here, we are given a 4-SAT instance in which there is an assignment where in at least fraction of the constraints, there are at least two literals that are set to be true. Under this promise, can we find an assignment in polynomial time where in at least $f(\epsilon)$ fraction of the constraints, at least one literal is set to be true, for some function with $f(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$? The analysis of [CMM09] for the 2-SAT problem where they compute the probability of a two dimensional Gaussian with a given mean and covariance matrix lying in the positive orthant, cannot be easily extended to dimensions. Instead, we follow a simpler analysis where we choose a single variable in the constraint carefully and show that with high probability, it gets rounded to True, and thus the predicate is satisfied. For a general predicate $(P; Q)$, we pick the coordinate by first showing that there is a halfspace separating the convex hull of using the fact that $(P; Q)$ contains Majority of all odd arities as polymorphisms. While our analysis is simpler and more general than [CMM09], we only achieve an error parameter $\epsilon^{1/3}$, similar to [Zwi98].

For the Alternating-Threshold (AT) case, we combine these ideas with a random geometric sampling trick. As a concrete example, consider the PCSP (3-SAT, NAE-3-SAT). For the exact case, we can solve the problem using basic SDP via hyperplane rounding: The vectors v_1, v_2, v_3 satisfy the property that the sum $v = v_1 + v_2 + v_3$ is along the direction of v_0 . Thus, we pick a hyperplane with a normal vector that is orthogonal to v_0 , and round x_i to be True if $\langle v_i, r \rangle > 0$, and False otherwise. For the robust setting, we get that the vector component normal to v_0 is small. Using this, we design a rounding scheme that is similar to the above, with the addition that when the vector's component normal to v_0 is small enough, we round it to True or False depending on its component along v_0 . The final ingredient is a geometric sampling trick where we sample the ratio of these two metrics randomly from a geometric series.

⁴Basic SDP (formally described in Section 6.3) is a well studied [Rag08] SDP relaxation of CSPs.

6.1.2 Unique Games based hardness

Unlike the algorithms part, in our hardness results, we crucially use the symmetry of the predicates. Furthermore, we assume that the PCSP contains a single predicate $(P; Q)$ that does not admit AT or MAJ polymorphisms, and we allow constraints to use negations of variables and unary constraints that set variables to be True or False. This is equivalent to asserting that the associated polymorphisms are folded and idempotent⁵. We show that for these Boolean symmetric folded idempotent PCSPs without AT and MAJ polymorphisms, the basic SDP relaxation has an integrality gap with perfect completeness, which by Raghavendra's framework connecting SDP gaps and Unique-Games hardness [Rag08], rules out robust satisfaction algorithms (under the Unique Games conjecture [Kho02a]).

More formally, we state our main hardness result below.

Theorem 64. For every Boolean symmetric folded idempotent PCSP $(P; Q)$ such that $AT_{L_1}; MAJ_{L_2} \not\subseteq \text{Pol}(\cdot)$ for some odd integers $L_1; L_2$, does not admit a robust algorithm assuming the Unique Games Conjecture.

As mentioned above, our Unique Games hardness (Theorem 64) is based on an integrality gap for the basic SDP relaxation. Towards this end, we present a general recipe for showing integrality gaps with respect to basic SDP for Promise CSPs via colorings of the n -dimensional unit sphere.

Consider a simple local rounding scheme for a PCSP $(P; Q)$: for every $n \geq 1$, there is a fixed rounding function $f_n : S^n \rightarrow \{-1, +1\}^g$. We consider the basic SDP relaxation and obtain $v_x \in \mathbb{R}^n$ corresponding to every variable x . We round it to an integral solution where we map v_x to $f_n(v_x)$. For this to be a valid rounding scheme, whenever a configuration of vectors $V = (v_1; v_2; \dots; v_k)$ can be assigned to the variables in a constraint, the corresponding integral values $(f_n(v_1); f_n(v_2); \dots; f_n(v_k))$ must belong to Q . Note that proving that there are no such local rounding functions is a necessary step towards showing integrality gaps for the PCSP using compactness arguments, we can show that this is sufficient as well.

As a concrete example, consider the CSP 3-LIN: a set of three vectors $v_1; v_2; v_3$ can be assigned to a set of vertices $x_2; x_3$ of a constraint by the basic SDP if the gram matrix of these vectors is in the convex hull of the gram matrices of the satisfying assignments. We refer to such a set of three vectors as a P-configuration. For the 3-LIN case, this condition can be translated to the fact that the three vectors are pairwise orthogonal. Thus, to show that basic SDP does not solve 3-LIN, it suffices to show that for some integer n , there is no function $f : S^n \rightarrow \{-1, +1\}^g$ such that whenever $v_1; v_2; v_3 \in S^n$ are mutually orthogonal, then there are odd number of ± 1 s in $(f(v_1); f(v_2); f(v_3))$. As we allow negation of variables, we also require such a function to be folded, i.e. $f(-v) = f(v)$ for every $v \in S^n$. In fact, it's easy to show such a coloring does not exist: consider a set of three mutually orthogonal vectors $V = (v_1; v_2; v_3)$ and their negations $V^0 = (-v_1; -v_2; -v_3)$. Note that both these are a valid P-configurations, but at least one of $(f(v_1); f(v_2); f(v_3))$, $(f(-v_1); f(-v_2); f(-v_3))$ must have an even number of ± 1 s, completing the proof that there is no such local rounding function. This corresponds to the textbook Basic SDP integrality gap instance 3-LIN consisting of the

⁵We say that a function $f : \{-1, +1\}^n \rightarrow \{-1, +1\}^g$ is folded iff $f(-x) = f(x)$. We say f is idempotent if $f(b; b; \dots; b) = b$ for every $b \in \{-1, +1\}^g$.

constraints $f(x_1, x_2, x_3); (\bar{x}_1, \bar{x}_2, \bar{x}_3)g$.

We can define the P -conjugations for an arbitrary PCSP $(P; Q)$, and use this approach to show the absence of sphere coloring “respecting” which in turn implies the required integrality gap with respect to the basic SDP for G -LIN. While the P -conjugations in the above proof for G -LIN are easy to study, in general, proving the absence of sphere coloring is challenging. For example, consider the PCSP $(P; Q)$ where $P = \{1; 5\}$ -in-5-SAT, $Q = \{1; 2; 3; 4; 5\}$ -in-5-SAT. Here, a set of P -conjugations are n unit vectors such that every two distinct vectors have inner product equal to $-\frac{1}{5}$. The sphere coloring problem is then to show that there exists n such that for any folding $f: S^n \rightarrow \{1, \dots, 5\}$, there exists a set of n vectors S^n with every pair of them having inner product equal to $-\frac{1}{5}$ that are all colored 1.

Such problems where the goal is to find a monochromatic structure in sphere colorings are studied under the title “sphere Ramsey theory”. In a striking result using tools from combinatorics, linear algebra, and Banach space theory, Meka and Rödl [MR95] proved that every set of n nearly independent vectors whose circumradius is smaller than $\frac{1}{\sqrt{n}}$ is sphere Ramsey—i.e., for every r , there exists n large enough such that every coloring of S^n must have a monochromatic set U that is congruent to U . This directly answers the above question regarding sphere coloring of $(P; Q)$ where $P = \{1; 5\}$ -in-5-SAT, $Q = \{1; 2; 3; 4; 5\}$ -in-5-SAT.

For an arbitrary Boolean symmetric PCSP $(P; Q)$, to prove Theorem 64, we first reduce the problem into a fixed number of templates using the properties of AT and MAJ polymorphisms [BG21b]. Then, we use the result of Meka and Rödl [MR95], and a connectivity lemma for conjugations to show the absence of sphere colorings for these templates, except for mentioned in Theorem 64.

Our results highlight a close connection between robust algorithms for PCSPs and being solved by the basic SDP. For a Promise PCSP $(P; Q)$, by being solved by the basic SDP, we mean that for every instance of a PCSP has an integral solution satisfying Q if and only if the basic SDP relaxation using P is feasible or ϵ . In other words, we can use the basic SDP to solve the decision version of the PCSP. As our algorithms for the AT and MAJ polymorphisms are based on rounding the basic SDP, we get that every Boolean folded PCSP that contains MAJ polymorphisms can be solved by the basic SDP. In the proof of Theorem 64, we showed that a vast majority of Boolean symmetric folded PCSPs with AT or MAJ polymorphisms cannot be solved by the basic SDP. This suggests a more general relation between the basic SDP and robust algorithms for PCSPs. Informally, for both the existence of robust algorithms and being solved by the basic SDP, the underlying requirement is the existence of polymorphism families that are robust to noise. While our results show that this is true for the PCSPs that we study in this chapter, we believe this is a more general phenomenon.

Conjecture 65. A PCSP has a polynomial time robust algorithm if and only if it can be solved by the basic SDP.

Note that if there is an integrality gap for with respect to the basic SDP relaxation, then by Raghavendra's [Rag08] result, we get that does not have a polynomial time robust algorithm, assuming the Unique Games Conjecture. This already proves one direction of Conjecture 65. The other direction is more interesting: can we obtain robust algorithms for PCSPs just using the

fact that basic SDP solves them exactly? We also remark that the conjecture is already proven to be true for CSPs, where the existence of robust algorithms [BK16] and solvability by basic SDP [TZ18] are both shown to be equivalent to having bounded width.

6.1.3 Minion characterization of basic SDP

In addition to our concrete characterization of robust algorithms for a subfamily of PCSPs, we also present a novel algebraic characterization of which PCSPs can be solved via basic SDPs. Originally, in the study of CSPs, such algebraic characterizations were structured as follows (e.g., [Bul17; Zhu20]).

- “Algorithm A solves CSP(), iff there is a polymorphism $\gamma \in \text{Pol}()$ with specific properties.”

A key property of the polymorphisms of a PCSPs that one can take minors by identifying coordinates. Relationships between a finite set of functions can be captured by identities, but more advanced relationships can be captured by infinite objects known as

Since the early days of PCSPs, it has been known that a single polymorphism cannot dictate hardness (c.f., [BG21b]), and thus one must instead consider a sequence of polymorphisms (e.g., [Bra+20]):

- “Algorithm A solves PCSP(), if and only if there is an infinite sequence of polymorphism $f_1; f_2; \dots \in \text{Pol}()$ with specific properties.”

However, in many cases, such a characterization is unfeasible or unwieldy. Instead a more general characterization, first characterized by [Bar+21], captures the structure of polymorphism via a minion (formally defined in Section 6.6).

- “Algorithm A solves PCSP(), if and only if there is minion homomorphism from \mathcal{M}_A to $\text{Pol}()$.”

Many recent papers [Bra+20; CZ22a; CZ22b] have proven such characterizations in various contexts. Our contribution to this line of work is showing that the basic SDP can be captured by a minion, which we call \mathcal{M}_{SDP} .

Theorem 66. The exact basic SDP solve PCSP() if and only if there is a minion homomorphism from \mathcal{M}_{SDP} to $\text{Pol}()$.

We note that the theorem applies equally to Boolean and non-Boolean PCSPs.

The construction of the \mathcal{M}_{SDP} minion is inspired by the vector interpretation of solutions to the Basic SDP. Each object in the minion is a collection of orthogonal vectors which sum to a reference vector r_0 . The minors involve adding groups of vectors together. Having a minion homomorphism from \mathcal{M}_{SDP} to $\text{Pol}()$ implies that there are polymorphisms whose minors behave exactly like combining orthogonal vectors.

Proving Theorem 66 has a few technical hurdles. One challenge is that SDP solutions may require vectors of an arbitrarily large dimension. In order for these arbitrarily-large dimensional relationships to be captured in our minion, we have that the families of vectors make \mathcal{M}_{SDP}

⁶Here ‘exact’ means that we verify that the basic SDP is solved to exact precision, within a bits of precision is the computational limit. A more thorough discussion of this technicality is in Section 6.6.

reside in a (countably) infinite-dimensional vector space. Similar techniques have been used in other minion constructions [Z22a; Z22b].

Another challenge that appears specifically unique to our work is that a Basic SDP solution gives a vector corresponding to each variable, but in order for the proof to go through additional vectors are needed which correspond to the clauses. [The variable vectors are "projections" of the clause vectors.] Obtaining such clause vectors would typically be done via Sum-of-Squares or a related routine, but we prove that including such clause vectors are without loss of generality. That is, for any Basic SDP solution, it can be extended to a solution which includes clause vectors without modifying the original Basic SDP solution. This gives us enough vector structure to prove that the minion homomorphism corresponds to the Basic SDP solution.

Path to integrality gaps. As a direct consequence of the minion structure theorem, we can connect sphere coloring with integrality gaps. By a result of [Bar+21], the minion homomorphism $M_{SDP} \rightarrow \text{Pol}(\cdot)$ is equivalent to finding a satisfying assignment to a "universal" instance of $\text{PCSP}(\cdot)$ known as a free structure. In the case that \cdot is a Boolean PCSP, this free structure for M_{SDP} turns out to be an instance where every possible unit vector is a variable. The clauses correspond to collections of vectors which satisfy the corresponding basic SDP. For the general theory of approximation of basic SDPs, similar constructs with sphere coloring being a "universal" gap have appeared in the literature ([Bra+21]). To make the integrality gaps more self-contained, we streamline the connection between sphere coloring and integrality gaps in Lemma 86.

Organization. We first start by introducing formal definitions and some general observations in Section 6.2 and Section 6.3. We provide our algorithmic results (Theorem 63) in Section 6.4 and prove the hardness results (Theorem 64) in Section 6.5. Finally, we study the basic SDP minion in Section 6.6.

6.2 Preliminaries

Notations. We use $[n]$ to denote the set $\{1, 2, \dots, n\}$. For k -ary relation $A \subseteq [q]^k$, we abuse the notation and use A both as a subset of $[q]^k$, and also as a predicate: $A: [q]^k \rightarrow \{0, 1\}$. For a vector $x = (x_1, x_2, \dots, x_n) \in \{0, 1\}^n$, we use $\text{hw}(x)$ to denote the number of 1s in x , i.e., $\text{hw}(x) = \sum_{i=1}^n x_i$. For $S \subseteq \{0, 1, \dots, k\}$, we use $\text{Ham}_k(S)$ to denote $\{x \in \{0, 1\}^k : \text{hw}(x) \in S\}$. We use NAE_k to denote the set $\text{Ham}_k(\{1, 2, \dots, k-1\})$. For vectors $x, y \in \mathbb{R}^n$, we use $x \cdot y$ and $\langle x, y \rangle$ interchangeably to denote $\sum_i x_i y_i$.

Boolean symmetric folded PCSPs We restrict ourselves to Boolean symmetric folded PCSPs in this chapter.

Definition 67 (Boolean symmetric folded PCSP) A PCSP $\langle (A_1; B_1); (A_2; B_2); \dots; (A_l; B_l) \rangle$ over a pair of domains D_1, D_2 is said to be Boolean symmetric folded if the following hold:

1. (Boolean) The domains D_1 and D_2 are both equal to $\{0, 1\}$.
2. (Symmetric) All the relations are symmetric i.e., for every $i \in [l]$, and x, y such that $\text{hw}(x) = \text{hw}(y)$, we have $x \in A_i$ if and only if $y \in A_i$, and similarly, $x \in B_i$ if and only if $y \in B_i$.

3. (Folded) We allow negating the variables i.e., there exists $\sigma \in \{0, 1\}^L$ such that $A_i = B_i = f(\sigma_i; +1); (+1; \sigma_i)g$.

AT and MAJ polymorphisms. We extensively study Alternate-Threshold (AT) and Majority (MAJ) polymorphisms in this chapter:

1. For an odd integer $L \geq 1$ and $x \in \{0, 1\}^L$, we have

$$AT_L(x) = \begin{cases} +1; & \text{if } x_1 + x_2 + x_3 + \dots + x_L > 0 \\ 1; & \text{otherwise.} \end{cases}$$

2. For an odd integer $L \geq 1$ and $x \in \{0, 1\}^L$, we have

$$MAJ_L(x) = \begin{cases} +1; & \text{if } x_1 + x_2 + \dots + x_L > 0 \\ 1; & \text{otherwise.} \end{cases}$$

We also use $AT_L(x_1; x_2; \dots; x_L)$ for $x_i \in \{0, 1\}^k$ (similarly for MAJ) when applying AT_L coordinatewise. For a predicate $P \in \{0, 1\}^k$, we use $AT_L(P)$ to denote the set $\{x_1; x_2; \dots; x_L \in P \mid AT_L(x_1; x_2; \dots; x_L) = 1\}$. We say that AT (and resp. MAJ) is in $\text{Pol}(\cdot)$ if AT_L (and resp. MAJ_L) is in $\text{Pol}(\cdot)$ for every odd integer $L \geq 1$. For a predicate $P \in \{0, 1\}^k$, we use $O_{AT}(P)$ (and similarly $O_{MAJ}(P)$) to denote the set $\{x \in P \mid AT_L(x) = 1\}$.

Relaxations of PCSPs. We say that a PCSP φ^0 is a relaxation of another PCSP φ if $\text{Pol}(\varphi^0) \supseteq \text{Pol}(\varphi)$. If φ^0 is a relaxation of φ , then there is a gadget reduction from φ^0 to φ . More formally, it is referred to as positive primitive promise reduction (ppp-reduction) from φ^0 to φ , or equivalently, as φ^0 is ppp-de nable from φ .

Definition 68 (ppp-de nability of PCSPs ([BG21b])). We say that a PCSP $\varphi^0 = (P^0, Q^0)$ containing a single pair of predicates of arity k is ppp-de nable from a PCSP over the same domain pair if there exists a fixed constant l and a PCSP instance using l (we also allow identifying variables together) over $k + l$ variables $x_1; x_2; \dots; x_k; y_1; y_2; \dots; y_l$ such that

1. If $(x_1; x_2; \dots; x_k) \in P^0$, then there exist $y_1; y_2; \dots; y_l$ such that $(x_1; x_2; \dots; x_k; y_1; y_2; \dots; y_l)$ satisfies the strong constraints in φ .
2. If there exists a satisfying assignment $(z_1; z_2; \dots; z_{k+l})$ to the weak constraints in φ , then $(z_1; z_2; \dots; z_k) \in Q^0$.

More generally, we say that φ^0 is ppp-de nable from φ if every predicate pair in φ^0 is ppp-de nable from φ . Brakensiek and Guruswami [BG21b] showed that if φ^0 is a relaxation of φ , then φ^0 is ppp-de nable from φ . As the ppp-reductions are polynomial time reductions, this shows that φ^0 can be reduced to φ in polynomial time. We show that the ppp-reductions also preserve the existence of robust algorithms.

Lemma 69. Suppose that the PCSP φ over a pair of domains $D_1; D_2$ is a relaxation of φ^0 over the same domain pair i.e. $\text{Pol}(\varphi) \supseteq \text{Pol}(\varphi^0)$. If φ has a polynomial time robust algorithm, then φ^0 has a polynomial time robust algorithm as well.

Proof. As proved in [BG21b], φ^0 can be ppp-reduced to φ . Given an instance φ^0 of φ^0 over a set of variables U , we output an instance φ containing $|U|$ original variables and a set of

dummy variables. For every constraint (P^0, Q^0) involving the variables u_1, u_2, \dots, u_k in I^0 , we have a set of dummy variables v_1, v_2, \dots, v_l and a set of constraints using $u_1, u_2, \dots, u_k, v_1, v_2, \dots, v_l$ as in Definition 68. Let $V^0 = U^0 \cup V$ denote the set of variables of I , where U^0 denotes the set of original variables corresponding to I^0 , and V are the set of dummy variables. We claim that this reduction preserves robust algorithms.

1. (Completeness). Suppose that there exists an assignment D_1 to I^0 satisfying all the constraints. Then, there is an assignment D_2 to the dummy variables which together with assigning D_1 to the original variables satisfies all the constraints in I .
2. (Soundness). Suppose that there is an assignment D_2 satisfying $1 - O(\epsilon)$ fraction of the constraints in I . As each dummy constraint set uses $O(1)$ constraints, we get that in at least $1 - O(\epsilon)$ constraint sets, all the constraints are satisfied. This shows that the assignment restricted to U^0 satisfies $1 - O(\epsilon)$ fraction of the constraints in I^0 .

□

Elementary properties of Gaussians. We prove a couple of elementary properties of Gaussian distribution that we use later. First, we prove the following anti-concentration inequality for the standard Gaussian random variable.

Proposition 70. Suppose that $X \sim N(0, 1)$ has the standard Gaussian distribution. Then, for every $\epsilon > 0$,

$$\Pr(|X| \leq \epsilon) = O(\epsilon^2).$$

Proof. We have

$$\Pr(|X| \leq \epsilon) = \int_{-\epsilon}^{\epsilon} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx = O(\epsilon^2).$$

□

We also state the following concentration inequality for 1-dimensional Gaussian. Proposition 71. Suppose that $X \sim N(0, \sigma^2)$ has Gaussian distribution with variance σ^2 . Then, for every $t > 0$,

$$\Pr(X \geq t) \leq e^{-\frac{t^2}{2\sigma^2}}.$$

6.3 General Observations

6.3.1 Basic SDP setup

For simplicity when analyzing SDP rounding, we shall fix $\{0, 1\}$ as our Boolean domain. If it helps, think of 1 as 'True' and 0 as 'False.' We shall use bold text for vectors in real space, but non-bold text for PCSP tuples.

For $x \in \{0, 1\}^k$, we let $v_x \in \mathbb{R}^{k+1}$ be the column vector whose first coordinate is 1 and the remaining k coordinates are x . We define the KZ vertex [KZ97] at x to be $M_x := v_x v_x^T$. For

Let $P = \{f \in \{0, 1\}^k\}$, we let $KZ(P)$ denote the convex closure of $\{M_x : x \in P\}$. For $\epsilon > 0$, let $KZ_\epsilon(P)$ denote the convex closure of $\{M_x : x \in P\}$ (that is you have error).

For an instance of a PCSP on n variables and m clauses, we let $(A_i; B_i)$ be the clause type for the i th clause and let $S_i \subseteq [n]$ be the subset of variables to which the clause is applied.

The Basic SDP⁸:

$$\begin{aligned} & \text{minimize: } \sum_{i=1}^m \lambda_i \\ & \text{subject to: } M \succeq 0 \\ & \lambda_i \in [0, 1]; \sum_i \lambda_i = 1 \\ & \lambda_i \in [0, 1]; M_{j_{S_i}, j_{S_i}} = 1 \\ & \lambda_i \in [0, 1]; M_{j_{S_i}, j_{S_i}} \in KZ_{\epsilon}(A_i); \end{aligned}$$

We say that basic SDP is feasible if the above objective function is zero. We say that the basic SDP solves the PCSP if for every $\epsilon > 0$ such that the basic SDP is feasible, there is an assignment from $\{0, 1\}^n$ to the variables that satisfies all the clauses with respect to the weaker constraints.

6.3.2 Generic RHS reduction to “notx”

Let F be a family of functions (e.g. MAJ of all odd arities). For $k \geq 1$, let Q_{nx} be shorthand for $\{f \in \{0, 1\}^k\}$.

Claim 72. An algorithm A is a robust algorithm for $\text{InvPol}(F)$ ⁹ if and only if for all $k \geq 1$, $P \subseteq \{0, 1\}^k$ and $x \in \{0, 1\}^n \in O_F(P)$, we have that A is a robust algorithm for $\text{PCSP}(P; Q_{nx})$:

Proof. The “only if” direction is trivial as $(P; Q_{nx}) \in \text{InvPol}(F)$.

For the “if” direction, consider an instance and replace every clause of the form $(A; B)$ with at most 2^k clauses of the form $(A; Q_{nx})$ all on the same set of variables. By assumption, A is a δ -vs $1 - \epsilon$ robust algorithm for these $(A; Q_{nx})$ clauses, where δ is the maximum of the tradeoff functions for all $(A; Q_{nx})$. It is easy to see that the assignment A produces is a δ -vs $1 - 2^k \epsilon$ algorithm for $\text{PCSP}(A; B)$, which is robust as δ depends only on the choice of template, which is independent of ϵ . \square

6.4 Robust Algorithms

6.4.1 CMM is a robust algorithm for MAJ

We show that the robust algorithm of Charikar, Makarychev, and Makarychev [CMM09] for 2-SAT generalizes to every PCSP that has Majority polymorphisms of all odd arities. First, we recall the algorithm.

⁷Technically S_i is a “subtuple” not a subset as order matters, but how this is handled will always be obvious.

⁸This formulation may not make it clear why this is in fact an SDP (i.e., the dependencies are linear), but this is the most compact way to say what is really going on.

⁹We define $\text{InvPol}(F)$ to be the set of finite promise templates with $F \in \text{Pol}(\cdot)$.

1. Given an instance containing n variables, solve the basic SDP and obtain the Gram matrix $M \in \mathbb{R}^{(n+1) \times (n+1)}$. Let α be the 0 th column of M minus the first entry. Let $M_{[n] \times [n]}$ be the lower-right $[n] \times [n]$ submatrix of M .
2. Sample an n -dimensional Gaussian $N(0; M_{[n] \times [n]})$. (Note that $M_{[n] \times [n]}$ is PSD.)
3. Set $x^0 = (\cdot)^{\frac{2}{3}}$.
4. For each $i \in [n]$, round as follows

$$x_i = \begin{cases} +1 & \text{if } x_i^0 \geq \alpha \\ -1 & \text{otherwise} \end{cases}$$

We shall prove the following analysis guarantee about our algorithm.

Theorem 73. Let $\text{MAJ} \in \text{Pol}(\cdot)$. Let \mathcal{C} be an instance of $\text{PCSP}(\cdot)$ for which there is a basic SDP solution with a completeness of α (i.e., the error value of the SDP solution is $1 - \alpha$, where m is the number of clauses). Then, the CMM algorithm above finds an 'integral' assignment to \mathcal{C} which satisfies $(1 - 3\epsilon)$ fraction of the clauses in expectation.

We analyze the algorithm clause by clause. Fix a clause using the predicate $(P; Q)$. Let k denote their arity i.e. $P; Q \in \{+1, -1\}^k$. Suppose that the basic SDP solution satisfies the constraint with probability α i.e., the local distribution is supported with (α) fraction of the predicates from \mathcal{P} . Our goal is to upper bound the probability that the rounded solution violates the constraint Q by a function of α and c . Using the fact that the expected value over all the constraints is at most α , we get our required robustness guarantee. More formally, we prove the following

Lemma 74. Let $(P; Q)$ be a clause in the instance. Presume the basic SDP gives a value of $1 - c$, the probability that Q is satisfied by the CMM algorithm is

$$O \left(\frac{1 - c}{\log \frac{1}{1 - c} + 2c} \right) \log \frac{1}{1 - c} + \frac{c}{A}$$

As the expected value over all the constraints is at most α and using Jensen's, we get that the net error probability is $O(1 - 3\epsilon)$, proving Theorem 73. Thus, it suffices to prove Lemma 74.

By Claim 72, it suffices to find a robust algorithm for $\text{PCSP}(P; Q_{n \times n})$ where $\epsilon \in \text{MAJ}(P)$.

Observe that CMM is "sign-symmetric" in the following sense: if a variable is replaced by its negation then the probabilities that variable are assigned are exactly interchange. Furthermore, replacing a variable with its negation in a clause $(P; Q)$ does not change whether $\text{MAJ} \in \text{Pol}(P; Q)$. Henceforth, we may assume that $Q = \{+1, -1\}^k$ (i.e., $x = (-1; 1; \dots; 1)$).

Let \mathcal{P} be the convex hull of \mathcal{P} , where the tuples are viewed as vectors. We prove the following property about \mathcal{P} using the fact that the $\text{PCSP}(P; Q)$ has Majority of all odd arities as polymorphisms.

¹¹We use O to denote a hidden constant which depends on the specific template

Lemma 75. Let $P = \{x \in \mathbb{R}^k : x_i \in \{0, 1\}, \sum_{i=1}^k x_i \geq \frac{2k-1}{3}\}$ be a predicate such that $\{0, 1\} \in \text{MAJ}(P)$. Then, there is a hyperplane separating P from the origin: there exists $w \in \mathbb{R}^k$, $w_i \geq 0$ and $\|w\|_1 = 1$ such that for every $x \in P$, $\sum_{i=1}^k w_i x_i \geq 0$.

Proof. Consider the following linear program,

$$\begin{aligned} & \text{maximize:} \\ & \text{subject to:} \quad \sum_{i=1}^k w_i = 1 \\ & \quad \quad \quad \sum_{i=1}^k a_i w_i \geq 0 \quad \forall a \in P \\ & \quad \quad \quad w_i \geq 0 \end{aligned}$$

It suffices to prove that the objective of this linear program is non-negative. To do this, we consider the dual program on variable $\lambda \in \mathbb{R}$ and $\alpha_a \in \mathbb{R}$ for $a \in P$:

$$\begin{aligned} & \text{minimize:} \\ & \text{subject to:} \quad \lambda \geq 0 \quad \text{(dual of)} \\ & \quad \quad \quad \sum_{a \in P} \alpha_a = 1 \\ & \quad \quad \quad \sum_{a \in P} \alpha_a a_i \geq 0 \quad \forall i \in [k] \quad \text{(dual of } w) \end{aligned}$$

As all the coefficients used in the LP are rational, we may assume that α_a are rational. Assume for sake of contradiction that there is a solution to the dual LP with $\lambda < 0$. Then, we have that for all $i \in [k]$,

$$\sum_{a \in P} \alpha_a a_i < 0$$

Let N be the least common denominator of the α_a . Consider the set of satisfying assignments to P where we take $2N \alpha_a$ copies of a for each $a \in P$. We also add an arbitrary element b to our set. As $\sum_{a \in P} \alpha_a a_i < 0$ for every $i \in [k]$, and $\sum_{a \in P} \alpha_a a_i N$ is an integer, we get that for every $i \in [k]$, $\sum_{a \in P} \alpha_a a_i N + b_i < 0$. In other words, when we apply MAJ to this set of assignments, we get $(x_1, x_2, \dots, x_k) \in P$. As $(P; Q)$ contains Majority of all odd arities as polymorphisms, this implies that the resulting output (x_1, x_2, \dots, x_k) is in Q , a contradiction.

Thus, the objective of the original LP is non-negative, completing the proof. \square

Now we use this lemma to complete the proof of Lemma 74. Suppose that the basic SDP solution satisfies the constraint $(P; Q)$ with error equal to ϵ i.e., there is a local distribution of $f \in \{0, 1\}^k$ that supports the vectors such that the weight of the assignments P is at most ϵ . Let $K = 2^k$. We have probabilities p_1, p_2, \dots, p_K corresponding to the local assignments $a_1, a_2, \dots, a_K \in \{0, 1\}^k$ where each $p_i \geq 0$ and $\sum_{i \in [K]} p_i = 1$ such that

$$\sum_{i \in [K]} p_i a_i = c$$

Using Lemma 75, we get $w \in \mathbb{R}^k$ with $w \geq 0$ and $\|w\|_1 = 1$ such that $w^T a_i \geq 0$ for all $a_i \in P$. Combining this with the above properties of the basic SDP solution, we get the following.

1. (First moment). We have

$$w^T = \sum_{i \in [K]} p_i w^T a_i$$

(Using Lemma 75 and $\sum_{i \in [K]} w^T a_i = 1$)

2. (Second moment). We have

$$w^T w = \sum_{i \in [K]} p_i (w^T a_i)^2 + c \sum_{i \in [K]} p_i w^T a_i + c w^T = \sum_{i \in [K]} p_i (w^T a_i)^2 + c w^T + 2c :$$

We do casework on the value w^T . First, consider the case that $w^T = \frac{q}{\log 1}$. As $\|w\|_1 = 1$, and $w \geq 0$, there exists $i \in [K]$ such that $w^T a_i \geq \frac{q}{N}$. As $w^T a_i \in (0, 1)$, using Proposition 71, with probability at least $1 - \epsilon$, we have $w^T a_i \geq \frac{q}{N} - \epsilon$. Thus, with probability at least $1 - \epsilon$, the rounded solution satisfies $e^{\frac{q}{N} - \epsilon}$.

Henceforth, we assume $w^T < \frac{q}{N}$. For notational convenience let $t = \frac{q}{N}$. We have

$$w^T t = \frac{c}{N} \tag{6.1}$$

and $w^T w = \sum_{i \in [K]} p_i (w^T a_i)^2 + 2c w^T + 2c$. Note that $w^T \in N(0, w^T w)$. Thus, using Proposition 71, with probability at least $1 - \epsilon$, we have that

$$|w^T - t| \leq O\left(\frac{r}{(\frac{q}{N} + 2c) \log \frac{1}{\epsilon}}\right) \tag{6.2}$$

Note that the rounded solution does not satisfy $w^T \geq t$ only if $w^T < t$. We now upper bound the probability that this can occur. Together with Equation (6.1) and Equation (6.2), implies that

$$0 \leq w^T - t \leq O\left(\frac{r}{(\frac{q}{N} + 2c) \log \frac{1}{\epsilon}} + \frac{c}{N}\right) :$$

Take some coordinate with $w_i = 1 = k$ and note that

$$t_i = \sum_{j \in [K]} p_j w_j^T a_j \leq O\left(\frac{r}{(\frac{q}{N} + 2c) \log \frac{1}{\epsilon}} + \frac{c}{N}\right) ;$$

but this can only happen with probability $O\left(\frac{p}{\frac{q}{N} + c \log \frac{1}{\epsilon}} + \frac{c}{N}\right)$ using Proposition 70. Thus, the probability that the rounded solution does not satisfy $w^T \geq t$ is at most

$$O\left(\frac{p}{\frac{q}{N} + c \log \frac{1}{\epsilon}} + \frac{c}{N}\right) :$$

This completes the proof of Lemma 74 and Theorem 73.

6.4.2 Warm-up for AT: Oblivious LP rounding algorithm for OR

As a stepping-stone for our algorithm for AT presented in the next section, we present a robust algorithm for the OR polymorphism. Horn-SAT is an example of a (P)CSP for which OR is a polymorphism. A robust algorithm for Horn-SAT was found previously by Zwick [Zwi98]. (See also the matching hardness result by Guruswami and Zhou [GZ12].) We now present an algorithm for PCSPs with the OR polymorphism achieving similar guarantees. As mentioned earlier, besides a polymorphic generalization of the Horn-SAT robust algorithm, our motivation is a warm-up for the algorithm for AT polymorphism in the next section.

As the OR operator is naturally over the $\{0, 1\}^k$ basis, we shall assume that the predicates $P; Q \in \{0, 1\}^k$ for this section.

1. Solve the basic LP and obtain the value y_i for each variable $i \in [n]$.
2. Let T be a geometric progression with first term 2^{-k} , last term $1 - 2^{-k}$ and spacing between terms is at least $\frac{1}{2^k}$, where k is an upper bound on the maximum clause size.
3. Sample a uniformly random threshold $t \in T$.
4. For each $i \in [n]$, round as follows

$$x_i = \begin{cases} 1 & \text{if } y_i \geq t \\ 0 & \text{otherwise} \end{cases}$$

Theorem 76. Let \mathcal{C} be a PCSP such that $\text{OR} \in \text{Pol}(\mathcal{C})$. Let $\mathcal{C}(P; Q)$ be an instance of PCSP(\mathcal{C}) for which there is a basic LP solution with a completeness of $\frac{1}{2}$. Then, with high probability our algorithm finds an integral assignment to which satisfies $1 - O(\frac{1}{2^k})$ fraction of the clauses in expectation.

Proof. Since we assume k is a constant, the size of T is $O(2^k)$. As with the analysis of MAJ, we fix a single clause $(P; Q)$ and analyze that. In fact, by Markov's inequality we may assume that $\frac{1}{2}$ fraction of the clauses have value at least $\frac{1}{2}$.

Since $\text{OR} \in \text{Pol}(P; Q)$, we in fact have that $\text{OR} \in \text{Pol}(Q)$ as well (see [BG21b]). Thus, by Schaefer's theorem [Sch76] PCSP(Q) and thus PCSP($P; Q$) can be pp \leq -reduced to Horn-SAT. Thus, we assume we are working with CNF clauses with at least $k - 1$ variables negated.

Consider a Horn-SAT clause on variables x_1, \dots, x_k with value at least $\frac{1}{2}$. First assume none of the variables are negated. Then, we have that $y_1 + \dots + y_k \geq \frac{1}{2}$. By pigeonhole, we must have some $(1 - y_i) \geq \frac{1}{2}$. Thus, $y_i \leq \frac{1}{2}$ for all $i \in [k]$. So the clause must be satisfied.

Otherwise, without loss of generality assume x_1 is negated. Then, we have that $(1 - y_1) + y_2 + \dots + y_k \geq \frac{1}{2}$. In particular, there exists $2 \leq i \leq k$ such that $y_i \geq \frac{1}{2}$. If $y_1 \geq \frac{1}{2}$, then $x_1 = 0$ with certainty. Otherwise, $y_i \geq \frac{1}{2}$, so there exists at most one $t \in T$ which would round y_i to 0 but y_1 to 1. Therefore, the probability the clause is satisfied is at least $1 - \frac{1}{|T|} = 1 - O(\frac{1}{2^k})$. This completes the proof. \square

6.4.3 Algorithm for AT

We now show how to combine ideas for MAJ and OR algorithms to give an algorithm for AT. Suppose that $\text{Pol}(\cdot)$ contains AT_L for every odd integer L . We state our algorithm below.

1. Solve the basic SDP and obtain vectors $v_1, \dots, v_n \in \mathbb{R}^n$.
2. Sample a vector $x \in \mathbb{R}^n$ by choosing each coordinate independently from $(0, 1)$.
3. Choose x uniformly at random from $\{p; r; \dots; r; p\}$ where $p = \frac{1}{2} - \epsilon$, and $r = \frac{1}{2} + \epsilon$ such that $p = \frac{\log \frac{1}{\epsilon}}{\log \log \frac{1}{\epsilon}}$.
4. For every $j \in [n]$, let $v_j = v_j + v_j^0$, where v_j^0 is orthogonal to v_j . We set x_j as follows.

$$x_j = \begin{cases} 1; & \text{if } |x_j - v_j| \leq |v_j^0| \\ +1; & \text{otherwise.} \end{cases}$$

We show that the above algorithm is a robust algorithm for every PCSP with polymorphisms.

Theorem 77. Let \mathcal{P} be a Boolean PCSP such that $\text{AT} \in \text{Pol}(\mathcal{P})$. Let \mathcal{I} be an instance of $\text{PCSP}(\mathcal{P})$ for which there is a basic SDP solution with completeness at least $\frac{1}{2}$. Then, the integral solution output by the above algorithm satisfies at least $\frac{1}{2} - \epsilon$ fraction of constraints of \mathcal{I} .

Theorem 77 together with Theorem 73 completes the proof of Theorem 63.

For ease of notation, we just use $O(\cdot)$ instead of $\Omega(\cdot)$ when ϵ is clear from the context. Consider an arbitrary predicate $P; Q$ with arity k i.e., $P; Q \in \{0, 1\}^k$. As $\text{AT} \in \text{Pol}(\mathcal{P}; Q)$ contains AT of all odd arities as polymorphisms, $\text{AT} \in \text{Pol}(\mathcal{P})$. Henceforth, in our analysis, we assume that $\text{AT} \in \text{Pol}(\mathcal{P})$. Furthermore, by using Markov's inequality, at least $\frac{1}{2}$ fraction of the constraints have SDP error at most ϵ . We restrict our analysis to these constraints with SDP error $\leq \epsilon$ using and and show that the rounded solution satisfies the predicate with probability at least $\frac{1}{2} - \epsilon$.

We study the case when \mathcal{P} and Q are symmetric in Section 6.4.3 and handle the general case in Section 6.4.4. We use the fact that \mathcal{P} and Q are symmetric to reduce to a special case. Note that if $Q = \{0, 1\}^k$, we are done. Henceforth, we assume that $Q \in \{0, 1\}^k$. Using the properties of symmetric PCSPs with AT polymorphisms proved by Brakensiek and Guruswami [BG21b], we get the following:

Lemma 78. Suppose that $(\mathcal{P}; Q)$ is a symmetric Boolean PCSP with arity k such that AT_L is a polymorphism of $(\mathcal{P}; Q)$ for every odd integer L , and $P; Q \in \{0, 1\}^k$, $Q \notin \{0, 1\}^k$. Then, either of the two following conditions hold:

1. $P = Q$ and $P \in \{0, 1\}^k; (+1; +1; \dots; +1)g$.
2. There exists $s \in \{0, 1, 2, \dots, k-1\}$ such that $P = \text{Ham}_k f; l_1; l_2g$, and $Q = \text{Ham}_k f; 1; 2; \dots; k-1g$.

Proof. Using Claim 4.6 in [BG21b], if $\text{Ham}_k f; l_1; l_2g \in P$ where $l_1 \notin l_2; f; l_1; l_2g \notin \{0, 1\}^k$, we get that $Q = \{0, 1\}^k$, a contradiction. Thus, either $P = \text{Ham}_k f; l; g$ for some $l \in \{0, 1, \dots, k\}$

or $P = \text{Ham}_k(f; 0; k)$. If $P \neq \text{Ham}_k(f; 0; k)$, we get that $P = Q$. If not, then $P = \text{Ham}_k(f; l; k)$ for $l \in \{1, 2, \dots, k-1\}$, and by the same Claim 6 in [BG21b], we get that $Q = O_{AT}(P) = \text{Ham}_k(f; 1; 2; \dots; k-1; k)$. \square

We consider the case where $P = Q$ and $P = \text{Ham}_k(f; 0; k)$.
 Lemma 79. For every constraint $(P; Q)$ where $P = Q$, $P = \text{Ham}_k(f; 0; k)$ such that the basic SDP has error at most ϵ on a constraint using $(P; Q)$, the above algorithm succeeds with probability $1 - O\left(\frac{\log \log \frac{1}{\epsilon}}{\log \frac{1}{\epsilon}}\right)$.

We defer the proof of Lemma 79 to Section 6.7.
 For the rest of the section, we assume that there exists $f; 1; 2; \dots; k-1; k$ such that $P = \text{Ham}_k(f; l; k)$, and $Q = \text{MAJ}_k$. More generally, we assume that there exists $R^k; b \in \mathbb{R}$ such that $w_i > 0 \forall i \in [k]$, $\sum_i w_i > b$, and $w_a = b$ for all $a \in P$. Here, for the case when $P = \text{Ham}_k(f; l; k)$, we can take $w = (1; 1; \dots; 1)$ and $b = 2^{l-k}$.

First, we prove some properties of the basic SDP vectors $v_1; \dots; v_k$ corresponding to the variables used in the constraint. As is the case with the MAJ algorithm, we have k probabilities $p_1; p_2; \dots; p_k$ corresponding to the assignments $a_1; \dots; a_k \in \{0, 1\}^k$ such that

$$\sum_{i \in [k]; a_i \in P} p_i = 1 - \epsilon$$

We use the basic SDP properties to get the following. Let $v_s = \sum_{i \in [k]} w_i v_i$, and let $v_s = v_0 + v_s^0$, where $\langle v_0, v_s^0 \rangle = 0$.

1. (First moments). We have

$$\sum_{i \in [k]} w_i \langle v_i, v_0 \rangle = \sum_{i \in [k]} w_i \langle v_i, v_s \rangle = \sum_{i \in [k]} p_i w_i a_i = b + \epsilon$$

where $\langle v_j, v_0 \rangle = O(\epsilon)$.

2. (Second moments). We have

$$\langle v_s, v_s \rangle = \sum_{i, j \in [k]} w_i w_j \langle v_i, v_j \rangle = \sum_{i \in [k]} p_i (a_i - w)^2 = b^2 + \epsilon$$

where $\langle v_j, v_0 \rangle = O(\epsilon)$.

Thus, we get $\langle v_s^0, v_s^0 \rangle = \langle v_s, v_s \rangle - (b + \epsilon)^2$ which is at most $O(\epsilon)$.

We are now ready to analyze the algorithm. We consider two cases separately:

Case 1. Suppose that there exists $l \in [k]$ such that $\langle v_l, v_s \rangle \leq \epsilon$. We claim that in this case, the rounded solution satisfies Q with probability at least $1 - O(\frac{1}{r})$.

Note that $v_j \in N(0; \frac{1}{2})$ for every $j \in [k]$. Suppose that we have $\langle v_j, v_s \rangle \leq \epsilon$ for some $j \in [k]$. Using Proposition 70, this implies that $\langle v_j, v_s^0 \rangle \leq \epsilon$ with probability at least $1 - \frac{1}{r}$. Furthermore, as $v_0 \in N(0; 1)$, using Proposition 71, we get that $\langle v_0, v_j \rangle \leq \epsilon$ with probability at least $1 - O(\frac{1}{r})$. Thus, with probability at least $1 - O(\frac{1}{r})$, x_j is set to be equal to 1 if $\langle v_j, v_s^0 \rangle > 0$, and 0 otherwise.

Hence, in order to show that the rounded solution satisfies it suffices to show that there exist $i_1, i_2 \in [k]$ such that $j_h; v_{i_1}^0$ and $j_h; v_{i_2}^0$ have opposite signs. As $kv_i^0 k_2 \leq kr^2$, with probability at least $O(\frac{1}{r})$, we have that $j_h; v_{i_1}^0$ and $j_h; v_{i_2}^0$ have opposite signs. Recall that $kv_s^0 k_2 \leq 0.25$. Thus, $j_h; v_{i_1}^0$ and $j_h; v_{i_2}^0$ have opposite signs with probability at least $O(\frac{1}{r})$. As $j_h; v_{i_1}^0$ and $j_h; v_{i_2}^0$ have opposite signs, there exists $i \in [k]$ such that $j_h; v_i^0$ and $j_h; v_{i_1}^0$ have opposite signs. Thus, with probability at least $O(\frac{1}{r})$, i and i_1 are rounded to different values, which implies that the rounded solution satisfies ϵ .

Case 2. Suppose that for every $i \in [k]$, we have $kv_i^0 k_2 \leq \frac{1}{2r^2}$.

As $j_h; v_i^0 \in N(0; kv_i^0 k_2)$, using Proposition 71, we get that with probability at least $O(\frac{1}{r})$, for every $i \in [k]$, $j_h; v_i^0 \leq \frac{1}{2r}$. On the other hand, using Proposition 70, we have that $v_{i_1}^0$ and $v_{i_2}^0$ have opposite signs with probability at least $O(\frac{1}{r})$. Furthermore, as $v_i^2 + kv_i^0 k_2^2 = 1$ for every $i \in [k]$, we get that $v_i^2 \geq 1 - \frac{1}{2r^2}$ for every $i \in [k]$. Thus, with probability at least $O(\frac{1}{r})$, for every $i \in [k]$, x_i is set to be ± 1 if $v_i^0 \geq 0$, and ∓ 1 otherwise. Combining this with the fact that $v_i^2 \geq 1 - \frac{1}{2r^2}$, and that $v_i^0 > b$ and $v_i^0 < -b$, for small enough r , we get the rounded solution has variables assigned ± 1 .

Completing the proof. We finish the proof by showing that with probability at least $O(\frac{1}{r})$, at least one of the above two cases hold. None of the above two cases hold if for some i we have

$$\frac{1}{2r^2} < kv_i^0 k_2 < kr^2$$

Or equivalently,

$$\frac{kv_i^0 k_2}{kr^2} < \frac{1}{2} < kv_i^0 k_2 2r^2$$

This holds with probability at most $O(\frac{1}{r})$ for every value of $kv_i^0 k_2$ as we are picking from $\{p; r; \dots; r\}$ uniformly at random.

6.4.4 General case for AT

For a vector $w \in \mathbb{R}^k$, define $\text{sgn}(w)_i$ to be ± 1 if $w_i \geq 0$ and ∓ 1 otherwise. Define AT to be the following family of weighted hyperplane predicates

$$\begin{aligned} \text{AT} &:= \{P_{w;b} := \sum_{i=1}^k w_i x_i = b; \\ Q_{w;b} &:= \sum_{i=1}^k \text{sgn}(w)_i x_i \geq b; \\ w &: \text{sgn}(w) > b; \quad w : \text{sgn}(w) < -b \} \end{aligned}$$

We observe that these predicates indeed have AT of all odd arities as polymorphisms.

Claim 80. $\text{AT} \in \text{Pol}(\text{AT})$

Proof. Fix $b \in \mathbb{Q}$ and $w \in (\mathbb{Q} \setminus \{0\})^k$. Let $(P_{w;b}; Q_{w;b})$ be the corresponding predicate for these values. It suffices to show that $\text{AT} \in \text{Pol}(P_{w;b}; Q_{w;b})$. Fix an odd arity L and pick $x_1; \dots; x_L \in \{P_{w;b}\}$. Observe that

$$\text{AT}(x_1; \dots; x_L) = \text{sgn}(x_1 + x_2 + \dots + x_L):$$

Further, $w(x_1 - x_2 + \dots + x_L) = b$. This implies that $\text{sgn}(x_1 - x_2 + \dots + x_L) \in \text{sgn}(w)$ as otherwise,

$$b = w(x_1 - x_2 + \dots + x_L) = w \cdot \text{sgn}(w) > b:$$

where we used the fact that the absolute value of each entry in $x_2 + \dots + x_L$ is at least 1. By a similar argument $\text{sgn}(x_1 - x_2 + \dots + x_L) \in \text{sgn}(w)$. Thus, $\text{AT}(x_1; \dots; x_L) \in Q_{w,b}$, as desired. \square

Note that our algorithm in the previous subsection gives a robust algorithm for these predicates as well. Henceforth, we reduce arbitrary Boolean PCSPs containing AT of all odd arities to these predicates via ppp reductions that preserve robust algorithms.

Let AT_{const} be the PCSP where constants can be specified. That is $\text{AT}_{\text{const}} = \{(f-1g); (f+1g); (f+1g); (f+1g)g\}$.
Theorem 81. Let P be any PCSP for which $\text{AT} \in \text{Pol}(P)$. Then, there is a ppp-reduction from AT_{const} to $\text{AT}_{\text{const}}[P]$.

Note that the analysis in Section 6.4.3 shows that our algorithm is a robust algorithm for AT_{const} . Furthermore, Lemma 79 shows that our algorithm is a robust algorithm for AT_{const} as well. Together with Theorem 81, we get that our algorithm is a robust algorithm for every PCSP that contains AT of all odd arities as polymorphisms.

To prove Theorem 81, we need to use the following lemma implicit in [BG21b]. We present the proof in Section 6.7 for the sake of completeness.

Lemma 82. Let P be a predicate such that there is non-trivial dependence in each coordinate (i.e., for each x_i , there exist assignments with $x_i = -1$ and $x_i = +1$). Then, $\text{O}_{\text{AT}}(P) = \{f \cdot \text{sgn}(x - y) : x, y \in A(P); \exists i, x_i \in \{-1, +1\}, y_i \in \{-1, +1\}\}$; where $A(P)$ is the affine hull of P .

Proof of Theorem 81 Fix a pair of predicates $(P; Q) \in \text{PCSP}$. It suffices to show that there is a PPP-reduction from $(P; Q)$ to $\text{AT}_{\text{const}}[P]$. If P has coordinates of fixed value, we can use a gadget reduction from AT_{const} to simulate these values. Thus, we assume that P has non-trivial dependence in each coordinate, and thus we apply Lemma 82 to get that $\text{O}_{\text{AT}}(P) = \{f \cdot \text{sgn}(x - y) : x, y \in A(P); \exists i, x_i \in \{-1, +1\}, y_i \in \{-1, +1\}\}$. We may without loss of generality assume that $Q = \{f \cdot \text{sgn}(x - y) : x, y \in A(P); \exists i, x_i \in \{-1, +1\}, y_i \in \{-1, +1\}\}$.

For every $x \in \{-1, +1\}^k \cap Q$, we find w, b such that $(P_{w,b} := \{f \cdot \text{sgn}(x - y) : w \cdot x = b\}; Q_{w,b} := \{f \cdot \text{sgn}(x - y) : w \cdot x = b, \text{sgn}(w) = x\})$ satisfy $P \subseteq P_{w,b}$ and $\text{sgn}(w) = x$. By applying this for every $x \in \{-1, +1\}^k \cap Q$, we get a set of predicate pairs $(P_1; Q_1); (P_2; Q_2); \dots; (P_L; Q_L)$ with $L \leq 2^k$ such that

1. $P \subseteq P_i$ for every $i \in [L]$.
2. $(P_i; Q_i) \in \text{AT}_{\text{const}}$ for every $i \in [L]$.
3. $\bigcap_{i \in [L]} Q_i = Q$.

This directly gives a ppp-reduction from $(P; Q)$ to $\text{AT}_{\text{const}}[P]$.

Henceforth, our goal is to show that for every $x \in \{-1, +1\}^k \cap Q$, we can find w, b such that $(P_{w,b} := \{f \cdot \text{sgn}(x - y) : w \cdot x = b\}; Q_{w,b} := \{f \cdot \text{sgn}(x - y) : w \cdot x = b, \text{sgn}(w) = x\})$ satisfy $P \subseteq P_{w,b}$, $w \cdot \text{sgn}(w) > b$, and $\text{sgn}(w) = x$. Without loss of generality, we can assume that $x = (+1, +1, \dots, +1)$. Fix an arbitrary vector $\bar{x} \in P$ such that $\bar{x} \geq f \cdot (-1, -1, \dots, -1); (+1, +1, \dots, +1)g$. Such a vector is

guaranteed to exist as does not contain \bar{x} and has non-trivial dependence on each coordinate. Let H be a subspace of \mathbb{R}^k defined as follows:

$$H := \{y \in \mathbb{R}^k : y \perp \bar{x}\}$$

As $\bar{x} \in O_{AT}(P)$, using Lemma 82, we get that for every $y \in H$, $\text{sgn}(z) \in x$, or in other words, there is no $y \in H$ with $z_i > 0$ for all $i \in [k]$. Using Claim 83, we can obtain w such that $w \cdot y = 0$ for all $y \in H$, and $w_i > 0$ for all $i \in [k]$. This shows that $w \cdot y = b$ for every $y \in P$, where $b = w \cdot \bar{x}$ satisfies $\sum_i w_i > b$, $\sum_i w_i > -b$. □

Claim 83. Let H be a subspace of \mathbb{R}^k such that there is no $y \in H$ with $y_i > 0$ for all i . Then, there exists w with $w_i > 0$ for all i and $w \cdot y = 0$ for all $y \in H$.

Proof. Since H and the positive orthant are both convex bodies, there exists $v \in \mathbb{R}^k$ and $b \in \mathbb{R}$ such that for all w in the positive orthant, $w \cdot v > b$ and for all $y \in H$, $y \cdot v \leq b$. Taking the limit as $w \rightarrow 0$, we have that $b \leq 0$. Further, since H is a subspace passing through the origin, we must have that $b = 0$ and that $y \cdot v = 0$ for all $y \in H$. Thus, v is normal to H . Note that v has all coordinates positive as $w \cdot v > 0$ for all w in the positive orthant. □

6.5 Unique Games based Hardness

In this section, we prove Theorem 64.

First, we use the analysis of AT and MAJ polymorphisms for symmetric PCSPs with folding and idempotence in [BG21b] to show that we can relax to one of five candidate PCSP types. Lemma 84. Let $\mathcal{P} = (P; Q)$ be a Boolean symmetric folded idempotent PCSP such that $\text{MAJ}_{L_1}; \text{AT}_{L_2} \geq \text{Pol}(\mathcal{P})$ for some odd integers L_1, L_2 . Then, there exists a PCSP $\mathcal{P}' = (P'; Q')$ that is a relaxation of \mathcal{P} that is equal to either of the following:

1. k is even, and $\mathcal{P}' = (P'; Q')$; $P' = \text{Ham}_k(f \frac{k}{2} g)$; $Q' = \text{Ham}_k(f 0; 1; \dots; k g)$ where $b \in \{f 1; k - 1 g\}$.
2. k is odd, $\mathcal{P}' = (P'; Q')$; $P' = \text{Ham}_k(f l; \frac{k+1}{2} g)$, $Q' = \text{Ham}_k(f 0; 1; 2; \dots; k - 1 g)$, where $l \in \{f \frac{k-1}{2} g\}$.
3. $\mathcal{P}' = (P'; Q')$; $P' = \text{Ham}_k(f l; k g)$, $Q' = \text{Ham}_k(f 1; 2; \dots; k g)$, where $l \in \{f 0; l - \frac{k-1}{2} g\}$.
4. $\mathcal{P}' = (P'; Q')$; $P' = \text{Ham}_k(f l g)$; $Q' = \text{Ham}_k(f 0; 1; \dots; k g)$ where $l \in \{f 0; k - 1 g\}$ where $l \in \{f 1; 2; \dots; k - 1 g; l - \frac{k-1}{2} g\}$.
5. $\mathcal{P}' = (P'; Q')$; $P' = \text{Ham}_k(f 1; k g)$; $Q' = \text{Ham}_k(f 0; 1; \dots; k g)$ for arbitrary $b \in \{f 0; 1; \dots; k g\}$.

We defer the proof of Lemma 84 to Section 6.7.

We show that the PCSPs $\mathcal{P}'_1, \mathcal{P}'_2, \mathcal{P}'_3, \mathcal{P}'_4$ and \mathcal{P}'_5 do not have robust algorithms by showing integrality gaps for the basic SDP relaxation of them. Raghavendra's result for CSPs [Rag08] shows that integrality gaps for the basic SDP relaxation can be translated to Unique Games Conjecture (UGC) [Kho02a] based inapproximability results. In fact, his result is verbatim applicable to Promise CSPs as well.

Theorem 85 (Special case of [Rag08] for PCSPs when the completeness value of SDP is 1). Suppose that for a Promise CSP, there is a finite integrality gap for the basic SDP i.e., there is a finite instance of Γ where basic SDP error is zero but is not satisfiable by Γ , even using the weak form of the constraints. Then, there exists a constant ϵ that is a function of Γ such that for every $\delta > 0$, assuming the Unique Games Conjecture, given an instance of Γ there is no polynomial time algorithm to distinguish between the two cases:

1. (Completeness) There exists an assignment that satisfies a fraction of the strong constraints.
2. (Soundness) No assignment satisfies a fraction of the weak constraints.

To obtain integrality gaps for the basic SDP relaxation for a PCSP, we study colorings of the n -dimensional sphere S^n that satisfies certain properties. First, we define certain notations that we need. For a predicate $P = \{ \pm 1 \}^k$, we say that a set of vectors $V = \{ v_1; v_2; \dots; v_k \}$ are a P -configuration with respect to another vector v_0 if the Gram matrix of $v_0; v_1; \dots; v_k$ is in $KZ(P)$. Fix a vector v_0 and we say that a coloring $f: S^n \rightarrow \{ \pm 1 \}^k$ respects the PCSP $(P; Q)$ if for every P -configuration $V = \{ v_1; v_2; \dots; v_k \}$ with respect to v_0 , we have that the colors of the vectors satisfy Q , i.e.,

$$(f(v_1); f(v_2); \dots; f(v_k)) \in Q$$

More generally, we say that a coloring $f: S^n \rightarrow D_2$ respects a PCSP over a pair of domains $D_1; D_2$ if it respects every predicate pair in Γ .

We show that the absence of such sphere coloring respecting Γ for some finite n gives an integrality gap for basic SDP relaxation of Γ .

Lemma 86. For every PCSP over a pair of domains $(D_1; D_2)$, the Basic SDP solves PCSP (Γ) if and only if for every $n \geq 1$, there exists a coloring $f: S^n \rightarrow D_2$ that respects Γ .

Proof. Via a compactness¹² argument (e.g., like the De Bruijn-Erdos theorem [BE51], for more details see Remark 7.13 of [Bar+21] or [Z22a]), we can infer that there is a coloring $f: S^n \rightarrow D_2$ respecting Γ if and only if for every finite subset $S \subseteq S^n$, there exists a coloring $f_S: S \rightarrow D_2$ that respects Γ .

First, assume that the Basic SDP solves PCSP (Γ) . For any finite subset $S \subseteq S^n$, we construct an instance of Γ where we add a constraint over S using the variables $x_{v_1}; \dots; x_{v_{r_i}}$ corresponding to the vectors $V = \{ v_1; v_2; \dots; v_{r_i} \}$ if V is a P_i -configuration. We have that $x_v \in D_1$ if v is a valid SDP solution. Thus, there exists an assignment to the variables that satisfies Γ , or equivalently, there exists $f_S: S \rightarrow D_2$ that respects the PCSP. And thus, there exists a coloring $f: S^n \rightarrow \{ \pm 1 \}^k$ that respects Γ .

Second, let $n \geq 1$ and assume that there exists a coloring $f: S^n \rightarrow D_2$ that respects Γ . We seek to show that the Basic SDP solves PCSP (Γ) . Take an arbitrary instance of PCSP (Γ) such that there is a solution to Basic SDP with objective value zero. Thus, there exists a dimension n and a mapping $\alpha: \Gamma \rightarrow \mathbb{R}^n$ of the variables to n -dimensional SDP vectors. By setting $\alpha(v_x) = f(v_x)$ for each variable x , we get an assignment satisfying all the constraints Γ . Thus, the Basic SDP solves Γ . \square

¹²We assume the axiom of choice.

Theorem 85 together with Lemma 86 shows that if a PCSP over a pair of domains $(D_1; D_2)$ does not admit a sphere coloring $S^n \rightarrow D_2$ that respects for some positive integer r , then, it does not admit a robust algorithm, assuming the Unique Games Conjecture. Thus, our goal is to show that the PCSPs mentioned in Lemma 84 do not admit sphere coloring that respects them. While we fail to achieve this for r_1 and r_4 , for the rest of the PCSPs, we show the absence of sphere coloring, which in turn implies Unique Games based hardness of obtaining robust algorithms.

In the rest of this section, we first prove a lemma regarding sphere Ramsey theory that we will use later. Then, we show that the earlier mentioned PCSPs do not have folded sphere coloring respecting them, thus showing that they don't admit robust algorithms. We remark that as we restrict ourselves to Boolean folded PCSPs, we only study the colorings $S^n \rightarrow \{-1, +1\}$ that are folded i.e. $f(-x) = f(x)$.

6.5.1 Sphere Ramsey Theory

For a finite set $S \subseteq \mathbb{R}^{n+1}$, we use $r(S)$ to denote the radius of the smallest sphere that contains S as a subset. Matser and Erdős [MR95] proved the following:

Theorem 87. Let S be the set of vertices of a simplex such that $\text{diam}(S) < 1$. Then, for every positive integer $r \geq 2$, there exists $n_0 := n_0(S; r)$ such that for every $n \geq n_0$, for every partition $f : S^n \rightarrow [r]$, there exists $S^0 \subseteq S^n$ that is monochromatic and is congruent to S .

We will use this to show the following lemma regarding sphere colorings.

Lemma 88. Fix an integer $k \geq 3$ and $r \geq 2$. There exists $n_0 := n_0(k)$ such that for every $n \geq n_0$ and coloring $f : S^n \rightarrow [r]$ and $2 \leq R < \frac{1}{k-1}$, there exists a monochromatic set of vectors $V = \{v_1; v_2; \dots; v_k\} \subseteq S^n$ such that $\|v_i - v_j\| = R$ for every $i \neq j$.

Proof. Consider an arbitrary set $S = \{u_1; u_2; \dots; u_k\}$ of k unit vectors in S^n such that $\|u_i - u_j\| = 1$ for every $i \neq j$. Such a set S is guaranteed to exist when n is large enough. We show that the vectors are nearly independent: suppose for contradiction that there exists scalars c_1, c_2, \dots, c_k not all zero, $\sum c_i = 0$ and $\sum c_i u_i = 0$. We have

$$0 = \sum_{i=1}^k c_i u_i = c_1 + (c_2 + \dots + c_k) = c_1 + (-c_1) = 0$$

implying that $c_1 = 0$. The same argument shows that $c_i = 0$ for all $i \in [k]$, a contradiction.

As S is nearly independent, they can be viewed as vertices of a simplex. Furthermore, the set of vectors can be embedded on a sphere of radius strictly smaller than R such that $0 < \epsilon < \frac{2}{k}$, and let $u_s = \sum_{i \in [k]} u_i, c = \|u_s\|$. We have

$$ku_s k_2^2 = \sum_i k u_i k_2^2 + 2 \sum_{i \in j} u_i \cdot u_j = k + \frac{k(k-1)}{2}$$

Note that

$$\begin{aligned} \| \sum u_i - c \|_2^2 &= \sum \|u_i - c\|_2^2 + k \|c\|_2^2 - 2c \cdot \sum u_i \\ &= 1 + \frac{k(k-1)}{2} \|c\|_2^2 - 2(1 + (k-1)c) \cdot c \\ &= 1 - k(1 + (k-1)\|c\|_2^2) \end{aligned}$$

which is strictly smaller than 1 when $0 < \|c\|_2 < \frac{2}{k}$. Thus, all the vectors are on a sphere centered at c and radius strictly smaller than 1, implying that $\delta(S) < 1$. Now, we can use Theorem 87 and f to obtain the required set of vectors. \square

While Theorem 87 is applicable to a wide range of sets, we sometimes need to apply it to sets S that do not form a simplex or have $\delta(S) = 1$. Towards this, we use the ‘‘Spreads’’ based idea in [MR95] to obtain a version of Theorem 87 directly for certain sets where Theorem 87 is not applicable.

We use the following notion of spread vectors from [MR95]. For an integer k , a vector $a \in \mathbb{R}^k$, and a set $J = \{j_1, j_2, \dots, j_k\}$ of cardinality k with $J \subseteq [n]$, we let

$$\text{Spread}_k(a; J) = \sum_{i=1}^k a_i e_{j_i}$$

where e_1, e_2, \dots, e_n is an orthonormal basis for \mathbb{R}^n . For a set $I \subseteq [n]$, we let

$$\text{Spread}_k(a; I) = \{ \text{Spread}_k(a; J) : J \subseteq I; |J| = k \}$$

We get the following as a direct application of the hypergraph Ramsey theorem.

Lemma 89. ([MR95]) For every $a \in \mathbb{R}^k; n; k$, there exists N such that in any coloring $\phi : \text{Spread}_k(a; [N]) \rightarrow [r]$, there exists J with $|J| = n$ such that $\text{Spread}_k(a; J)$ is monochromatic with respect to ϕ , i.e., $\exists p \in [r]$ such that $\phi(v) = p$ for all $v \in \text{Spread}_k(a; J)$.

Lemma 89 implies the following immediately.

Corollary 90. Suppose that $U = \{u_1, u_2, \dots, u_k\}$ be a set of k unit vectors such that $\|u_i - u_j\|_2 \geq \delta$ for $i \neq j$ and $\delta \in (0, \frac{1}{k}]$ for an integer N , and a vector $a \in \mathbb{R}^N$ with $\|a\|_2 = 1$. Then there exists $n_0 := n_0(U; a; N; \delta)$ such that for every $n \geq n_0; r$, for every sphere coloring $\phi : S^n \rightarrow [r]$, there exists a set of k vectors $V = \{v_1, \dots, v_k\}$ that are all colored the same, and $v_j = u_i - u_j$ for every $i; j \in [k]$.

We use Corollary 90 to obtain a couple of lemmas regarding sphere colorings. For ease of notation, we call a set of unit vectors $V = \{v_1, v_2, \dots, v_k\}$ to be k -regular if $\|v_i - v_j\|_2 = \frac{1}{k-1}$ for every $i \neq j$.

Lemma 91. Fix an integer $k \geq 2$. There exists $n_0 := n_0(k)$ such that for every $n \geq n_0$ and folded coloring $f : S^n \rightarrow [k-1; +1]$, there exist a k -regular set of vectors $V = \{v_1, v_2, \dots, v_k\} \subseteq S^n$ such that exactly $k-1$ vectors in V are colored -1 .

Proof. We construct a set of unit vectors $V = \{v_1, v_2, \dots, v_k\}$ in $\text{Spread}_N(a; [N])$ such that $\{v_1, v_2, \dots, v_{k-1}, v_k\}$ is a k -regular set, where N, a depend only on k , and $k \cdot a = 1$. Using Corollary 90, we can infer that in the coloring there exist k vectors $V = \{v_1, v_2, \dots, v_k\}$ that are all assigned the same color, such that $\{v_1, v_2, \dots, v_{k-1}, v_k\}$ is a k -regular set. As is folded, this implies that there is a regular set in which exactly $k-1$ vectors are assigned the color 1.

Thus our goal is to construct unit vectors $V = \{v_1, v_2, \dots, v_k\}$ in $\text{Spread}_N(a; [N])$ such that $\{v_1, v_2, \dots, v_{k-1}, v_k\}$ is a k -regular set. Or equivalently, we construct the vectors $v_1, v_2, \dots, v_{k-1} \in \text{Spread}_N(a; [N])$ and $v_k \in \text{Spread}_N(a; [N])$ such that $\{v_1, \dots, v_k\}$ is a k -regular set. We set $a = \frac{1}{2(k-1)}$ and $a = (\dots; \dots; \dots; \dots) \in \mathbb{R}^{2(k-1)}$. We set $v_i = \text{Spread}_N(a; J_i); i \in [k-1]; v_k = \text{Spread}_N(a; J_k)$ where J_1, J_2, \dots, J_k such that $|J_i| = 2(k-1)$ for every $i \in [k]$. We obtain these sets by induction. First, we consider the base case when $k=2$. In this case, we set $J_1 = J_2 = \{1, 2\}$ and $N = 2$ suffices. The vectors are the following:

$$\begin{aligned} v_1 &= (\dots; \dots) \\ v_2 &= (\dots; \dots) \end{aligned}$$

where $a = \frac{1}{2}$. Note that the above two vectors are a 2-regular set, and $v_1 \in \text{Spread}_2(a; [2]); v_2 \in \text{Spread}_2(a; [2])$ with $a = (\dots; \dots)$. Now, suppose that $v_1, v_2, \dots, v_k; N$ are such that $v_i = \text{Spread}_N(a; J_i); i \in [k-1]; v_k = \text{Spread}_N(a; J_k)$ satisfy the property that $\{v_1, v_2, \dots, v_k\}$ is a k -regular set with $a = (\dots; \dots; \dots; \dots) \in \mathbb{R}^{2(k-1)}$, $a = \frac{1}{2(k-1)}$. We construct $J_1^0, J_2^0, \dots, J_{k+1}^0$ such that $v_i^0 = \text{Spread}_{N^0}(a^0; J_i^0); i \in [k]; v_{k+1}^0 = \text{Spread}_{N^0}(a^0; J_{k+1}^0)$ satisfy the property that $\{v_1^0, v_2^0, \dots, v_{k+1}^0\}$ is a $(k+1)$ -regular set with $a^0 = (\dots^0; \dots^0; \dots^0; \dots^0) \in \mathbb{R}^{2k}$, $a^0 = \frac{1}{2k}$.

1. For every $i \in [k-1]$, we obtain J_i^0 from J_i by adding two new elements.

$$J_i^0 = J_i \cup \{N+2i; N+2i+1\}$$

This ensures that $v_i^0 = (0)^2$ for every $i; j \in [k-1]; i \notin j$.

2. We obtain J_{k+1}^0 from J_k by adding two new elements.

$$J_{k+1}^0 = J_k \cup \{N+1; N+2k\}$$

This ensures that $v_{k+1}^0 = (0)^2$ for every $i \in [k-1]$.

3. Finally, we set J_k .

$$J_k = \{N+1; N+2; \dots; N+2k\}$$

This ensures that $v_k^0 = (0)^2$ for every $i \in [k+1]; i \notin k$.

We illustrate our construction by obtaining the vectors for the case when $k=4$ and $k=4$:

$$\begin{aligned} v_1 &= (\dots; \dots; \dots; \dots; 0; \dots; \dots; \dots; 0) \\ v_2 &= (\dots; 0; 0; \dots; \dots; \dots; \dots; \dots; \dots) \\ v_3 &= (\dots; \dots; \dots; \dots; \dots; \dots; 0; \dots; \dots) \end{aligned}$$

$$\begin{aligned}
v_1 &= (\quad ; \quad ; \quad 0; \quad ; \quad ; \quad 0; \quad ; \quad ; \quad 0; \quad 0; \quad 0) \\
v_2 &= (\quad 0; \quad 0; \quad ; \quad ; \quad ; \quad ; \quad 0; \quad 0; \quad ; \quad ; \quad 0) \\
v_3 &= (\quad 0; \quad 0; \quad 0; \quad 0; \quad 0; \quad ; \quad ; \quad ; \quad ; \quad ; \quad) \\
v_4 &= (\quad ; \quad ; \quad ; \quad 0; \quad 0; \quad ; \quad ; \quad 0; \quad 0; \quad 0; \quad)
\end{aligned}$$

where $\alpha = \frac{1}{2}$ and $\beta = \frac{1}{6}$.

As the pairwise inner product of every pair (v_i^0, v_j^0) is equal to $(\alpha - \beta)^2 = \frac{1}{k}$, we get that these set of vectors are $(k-1)$ -regular set, completing the inductive proof. \square

6.5.2 Absence of sphere coloring

First, we show the absence of sphere coloring respecting using Lemma 91.

Lemma 92. Fix an even integer $k \geq 4$. There exists an integer n_0 such that for every $n \geq n_0$, there is no folded $f : S^n \rightarrow \{-1, +1\}$ that respects $\mu_1 = (P; Q); P = \text{Ham}_k f; \frac{k}{2}g; Q = \text{Ham}_k f; 0; 1; \dots; k-1g$ where $b \geq 1; k \geq 1g$.

Proof. Suppose for contradiction that such a folded exists. Fix a vector $v_0 \in S^n$. We get the P-conjugation of vectors $v_1; v_2; \dots; v_k$ such that the gram matrix of $v_0; v_1; \dots; v_k$ is a uniform convex combination from the vertices $\mathcal{KZ}(P)$. The vectors satisfy this if we have

1. (First moments) $\langle v_i, v_0 \rangle = 0$ for every $i \in [k]$.
2. (Second moments) $\langle v_i, v_j \rangle = \frac{2 \binom{k}{2} \frac{k^2}{4}}{\binom{k}{2}} = \frac{1}{k-1}$.

We restrict ourselves to vectors orthogonal to v_0 and apply Lemma 91 to obtain a set of vectors $v_1; v_2; \dots; v_k$ such that $\langle v_i, v_j \rangle = \frac{1}{k-1}$ and exactly $k-1$ of $f(v_1; v_2; \dots; v_k)$ are colored -1 . By negating these vectors if needed, we get a set of vectors $v_1; v_2; \dots; v_k$ such that $\langle v_i, v_j \rangle = \frac{1}{k-1}$ and exactly 1 of $f(v_1; v_2; \dots; v_k)$ are colored -1 . Thus, there exists a P-conjugation of vectors whose value contains exactly $k-1$ -1 s, a contradiction. \square

We show the absence of sphere coloring respecting μ_3 and μ_4 using Lemma 88.

Lemma 93. Fix an odd integer $k \geq 3$ and integer $s : 0 \leq s \leq \frac{k-1}{2}$. There exists an integer n_0 such that for every $n \geq n_0$, there is no folded $f : S^n \rightarrow \{-1, +1\}$ that respects $\mu_2 = (P; Q); P = \text{Ham}_k f; \frac{k+1}{2}g; Q = \text{Ham}_k f; 0; 1; 2; \dots; k-1g$.

Proof. Fix $v_0 \in S^n$. The P-conjugation that we consider is a set of vectors $v_1; v_2; \dots; v_k$ such that the gram matrix of $v_0; v_1; \dots; v_k$ is obtained by first sampling $i \in [k]$ where we set

$$i = \begin{cases} 1 & \text{with probability } \frac{1}{s} \\ \frac{k+1}{2} & \text{with probability } \frac{s}{s} \end{cases}$$

where $s = 1 - (k-1) < 0$. Then, we sample a uniformly random vertex $v \in \mathcal{KZ}(\text{Ham}_k f; g)$. We obtain the following properties:

1. (First moments) $v_i \cdot v_0 = 0$ for every $i \in [k]$.
2. (Second moments) $v_i \cdot v_j = \frac{1}{k}$ for every $i \neq j$. Furthermore, we get that $\frac{1}{k} < 1$.

Now, restricting ourselves to the vectors S^n that are orthogonal to v_0 , and using Theorem 87, we get that there exists a P -con guration of vectors that are all colored the same. By taking the negation of these vectors if needed, we get our required claim. \square

Lemma 94. Fix integers k, l such that $0 < l \leq \frac{k-1}{2}$. Then, there exists an integer n_0 such that for every $n \geq n_0$, there is no folded $f : S^n \rightarrow \{1, \dots, k\}$ that respects $\mathcal{P}_3 = (P; Q); P = \text{Ham}_k f l; k g$, $Q = \text{Ham}_k f l; 2; \dots; k g$, where $l \in [0; \frac{k-1}{2}]$.

Proof. Fix $v_0 \in S^n$. We pick the P -con guration along the same lines as in Lemma 93. We sample $i \in [k]$ with

$$i = \begin{cases} l & \text{with probability } \frac{k-l}{k-s} \\ k & \text{with probability } \frac{s}{k-1} \end{cases}$$

where $s = l - (k-l) < 0$. As before, we sample a uniformly random vertex $x \in \text{KZ}(\text{Ham}_k f i g)$. We get

1. (First moments) $v_i \cdot v_0 = \frac{k-l}{k-s} \frac{s}{k} + \frac{s}{k-1} \cdot 1 = 0$ for every $i \in [k]$.
2. (Second moments). For every $i, j \in [k]$, we get

$$v_i \cdot v_j = \frac{k-l}{k-s} \left(\frac{1}{2} + \frac{k-l}{2} \frac{l(k-l)}{k} \right) + \frac{s}{k-1}$$

The function $\frac{1}{2} + \frac{k-l}{2} \frac{l(k-l)}{k}$ is a decreasing function when $l \in [1; \frac{k-1}{2}]$. When $l = \frac{k-1}{2}$, the value of it is equal to $\frac{1}{k}$. Thus, we get that for every $l \in [1; \frac{k-1}{2}]$, we get that $\frac{1}{2} + \frac{k-l}{2} \frac{l(k-l)}{k} \geq \frac{1}{k}$. Thus, we get that

$$v_i \cdot v_j = \frac{1}{k}$$

for every $i \neq j$, and $\frac{1}{k} < 1$. We restrict ourselves to vectors S^n that are orthogonal to v_0 , and applying Theorem 87, we get that for any coloring $f : S^n \rightarrow \{1, \dots, k\}$, there is a monochromatic P -con guration that we described. By negating the vectors if needed, we get our required proof. \square

Lemma 95. Fix integers $k \geq 3; l \in [1; \dots; k-1]; l \leq \frac{k-1}{2}$. There exists integer n_0 such that for every $n \geq n_0$, there does not exist coloring $f : S^n \rightarrow \{0, 1\}$ that is folded i.e. $f(-x) = f(x)$, and respects the PCSP $\mathcal{P}_4 = (P; Q); P = \text{Ham}_k f l g; Q = \text{Ham}_k f 0; 1; \dots; k n f 0; k-1 g$.

Proof. We partition the predicate \mathcal{P} into P_1 and P_{-1} depending on the value of the first element, i.e.,

$$P_i = \{x \in P : x_1 = i; i \in \{1, -1\}\}$$

We pick the P -con guration as follows: sample $i \in \{1, -1\}$ uniformly at random, then, sample a uniformly random vertex $x \in \text{KZ}(P_i)$. We get

1. (First moments). By our choice of v_i s, we get that

$$v_1 - v_0 = 0$$

By using symmetry of the rest of the variables, we get that

$$v_i - v_0 = \frac{2l - k}{k - 1} \delta_{i \in \{2, 3, \dots, k\}}$$

For ease of notation, let $\alpha = \frac{2l - k}{k - 1}$.

2. (Second moments). We get

$$v_1 - v_i = \frac{1}{k - 1} \delta_{i \in \{2, 3, \dots, k\}}$$

$$v_i - v_j = \frac{(2l - k)^2 - (k - 2)}{(k - 1)(k - 2)} \delta_{i, j \in \{2, 3, \dots, k\}; i \neq j}$$

For ease of notation, let $\beta = \frac{(2l - k)^2 - (k - 2)}{(k - 1)(k - 2)}$.

We collect the following fact for later use:

$$\begin{aligned} \sum_{i=1}^k v_i^2 &= k + 2 \sum_{i, j \in \{2, 3, \dots, k\}; i \neq j} (v_i - v_j)^2 \\ &= k + 2 \sum_{i \in \{2, 3, \dots, k\}} (v_i - v_0)^2 + \sum_{i, j \in \{2, 3, \dots, k\}; i \neq j} (v_i - v_j)^2 \\ &= k + 2(k - 1) \frac{1}{k - 1} + (k - 1)(k - 2) \frac{(2l - k)^2 - (k - 2)}{(k - 1)(k - 2)} \\ &= (2l - k)^2. \end{aligned}$$

Thus, we get

$$\sum_{i=1}^k v_i = k - 2l \tag{6.3}$$

Our goal is to show that there exists ϵ such that for every $n \geq n_0$, for every folded sphere coloring $f : S^n \rightarrow \{-1, +1\}^k$ and $v_0 \in S^n$, there exists a set of vectors $V = \{v_1, v_2, \dots, v_k\}$ that satisfy the above first and second moments, and exactly ϵ vectors in V are colored $+1$, where $\epsilon \geq \frac{1}{2} - \frac{1}{k}$. For $i \in [k]$, let v_i^0 be the component of v_i orthogonal to v_0 :

$$v_i^0 = v_i - (v_i \cdot v_0) v_0; i \in [k]$$

Note that $\|v_i^0\| = \sqrt{1 - (v_i \cdot v_0)^2}$ for $i \in \{2, 3, \dots, k\}$ and $\|v_1^0\| = 1$. We let $u_i = \frac{v_i^0}{\|v_i^0\|}$. We have

$$\begin{aligned} u_i \cdot u_j &= \frac{v_i^0 \cdot v_j^0}{\|v_i^0\| \|v_j^0\|} \\ &= \frac{(v_i - v_0) \cdot (v_j - v_0)}{\sqrt{1 - (v_i \cdot v_0)^2} \sqrt{1 - (v_j \cdot v_0)^2}} \\ &= \frac{1}{2} \delta_{i, j \in \{2, 3, \dots, k\}; i \neq j} \end{aligned}$$

For ease of notation, $\rho = \frac{1}{1-\frac{2}{k}}$.

Using Equation (6.3), we get the following bound that we will use later.

We have

$$\begin{aligned} \|k-2\| &= \sum_{i=1}^k \|v_i\| \\ &= \|v_1\| + \sum_{i=2}^k \|v_i\| + (k-1)\|v_0\| \\ &= \|u_1\| + \rho \sum_{i=2}^k \frac{1}{1-\frac{2}{k}} \|u_i\| + (2\|v_0\| - k)\|v_0\| \end{aligned}$$

As $\langle u_i, v_0 \rangle = 0$ for every $i \in [k]$, we get that

$$\|u_1\| + \rho \sum_{i=2}^k \frac{1}{1-\frac{2}{k}} \|u_i\| = 0$$

Thus,

$$\sum_{i=2}^k \frac{1}{1-\frac{2}{k}} \|u_i\|^2 = \frac{1}{1-\frac{2}{k}}$$

On the other hand,

$$\begin{aligned} \sum_{i=2}^k \frac{1}{1-\frac{2}{k}} \|u_i\|^2 &= \sum_{i,j \in \{2, \dots, k\}; i \neq j} \frac{1}{1-\frac{2}{k}} \langle u_i, u_j \rangle \\ &= \sum_{i,j \in \{2, \dots, k\}} \frac{1}{1-\frac{2}{k}} \langle u_i, u_j \rangle \end{aligned}$$

Thus, we get

$$\sum_{i,j \in \{2, \dots, k\}} \frac{1}{1-\frac{2}{k}} \langle u_i, u_j \rangle = \frac{1}{(1-\frac{2}{k})(k-1)(k-2)} \quad (6.4)$$

We apply Lemma 88 on the following coloring of the sphere. For a vector $u \in S^n$ such that $\langle u, v_0 \rangle = 0$, let $f^0: S^{n-1} \rightarrow \{1, -1\}$ be defined as

$$f^0(u) = \begin{cases} 1 & \text{if } \langle u, v_0 \rangle + \rho \sum_{i=2}^k \frac{1}{1-\frac{2}{k}} \langle u, u_i \rangle \geq 0 \\ -1 & \text{otherwise} \end{cases}$$

Using Lemma 88 of [1] combined with the fact that $\sum_{i=2}^k \frac{1}{1-\frac{2}{k}} \langle u_i, u_j \rangle = \frac{1}{(1-\frac{2}{k})(k-1)(k-2)}$ obtained from Equation (6.4), we can infer that there exist $k-1$ unit vectors $u_1, u_2, \dots, u_{k-1} \in S^n$ such that $\langle u_i, v_0 \rangle = 0$ for all i , $\langle u_i, u_j \rangle = \frac{1}{(k-1)(k-2)}$ for all $i \neq j$ and $f^0(u_i) = f^0(u_j)$ for all $i \neq j$; $i, j \in [k-1]$.

We define $v_1^{(1)}, v_2^{(1)}, \dots, v_k^{(1)}; v_1^{(2)}, \dots, v_k^{(2)}$ as follows. For $i \in \{2, 3, \dots, k\}$, we let

$$\begin{aligned} v_i^{(1)} &= v_0 + \rho \sum_{j=2}^k \frac{1}{1-\frac{2}{k}} \langle u_j, v_0 \rangle v_j \\ v_i^{(2)} &= v_0 - \rho \sum_{j=2}^k \frac{1}{1-\frac{2}{k}} \langle u_j, v_0 \rangle v_j \end{aligned}$$

We let

$$v_1^{(1)} = \frac{\sum_{i=1}^{k-1} u_i}{\sum_{i=1}^{k-1} u_i}$$

and $v_1^{(2)} = v_1^{(1)}$. We now prove that the set of vectors $v_1^{(1)}; v_2^{(1)}; \dots; v_k^{(1)}$ and the set of vectors $v_1^{(2)}; v_2^{(2)}; \dots; v_k^{(2)}$ are a P -conjugation with first and second moments as computed earlier, where we sampled from $f^{-1}; +1g$ uniformly at random and sampled a uniformly random vertex from $KZ(P_i)$.

1. (First moments). As $v_0 = 0$ for all $i \in [k-1]$, we get that

$$v_1^{(1)} v_0 = v_1^{(2)} v_0 = 0$$

and

$$v_i^{(1)} v_0 = v_i^{(2)} v_0 = \frac{1}{k-1} \sum_{j=2}^k u_j$$

2. (Second moments). We have

$$\begin{aligned} v_1^{(1)} v_i^{(1)} &= \frac{\left(\sum_{j=1}^{k-1} u_j\right) \left(v_0 + \frac{1}{k-1} \sum_{j=1}^{k-1} u_j\right)}{\sum_{j=1}^{k-1} u_j} \\ &= \frac{1}{k-1} \frac{\left(\sum_{j=1}^{k-1} u_j\right) \left(\sum_{j=1}^{k-1} u_j\right)}{\sum_{j=1}^{k-1} u_j} \\ &= \frac{1}{k-1} \sum_{j=1}^{k-1} u_j \\ &= \frac{1}{k-1} \sum_{j=1}^{k-1} \frac{1}{k-1+2\frac{(k-1)(k-2)}{2}} \\ &= \frac{1}{k-1} \frac{1}{1-\frac{1}{k-1}} \text{ Using Equation (6.4)} \\ &= \frac{1}{k-1} \sum_{j=2}^k u_j \end{aligned}$$

Furthermore,

$$\begin{aligned} v_i^{(1)} v_j^{(1)} &= \left(v_0 + \frac{1}{k-1} \sum_{j=1}^{k-1} u_j\right) \left(v_0 + \frac{1}{k-1} \sum_{j=1}^{k-1} u_j\right) \\ &= \frac{1}{k-1} \sum_{j=2}^k u_j \\ &= \frac{1}{k-1} \sum_{j=2}^k u_j; i \neq j \end{aligned}$$

Similarly, we have

$$\begin{aligned} v_1^{(2)} v_i^{(2)} &= \frac{\left(\sum_{j=1}^{k-1} u_j\right) \left(v_0 + \frac{1}{k-1} \sum_{j=1}^{k-1} u_j\right)}{\sum_{j=1}^{k-1} u_j} = \frac{1}{k-1} \sum_{j=2}^k u_j \\ v_i^{(2)} v_j^{(2)} &= \left(v_0 + \frac{1}{k-1} \sum_{j=1}^{k-1} u_j\right) \left(v_0 + \frac{1}{k-1} \sum_{j=1}^{k-1} u_j\right) \\ &= \frac{1}{k-1} \sum_{j=2}^k u_j = \frac{1}{k-1} \sum_{j=2}^k u_j; i \neq j \end{aligned}$$

Thus, the set of vectors $v_1^{(1)}; v_2^{(1)}; \dots; v_k^{(1)}$ and the set of vectors $v_1^{(2)}; v_2^{(2)}; \dots; v_k^{(2)}$ are a P-conjugation. As $f(v_1^{(1)}) = f(v_2^{(1)}) = \dots = f(v_k^{(1)})$, we can infer that $f(v_1^{(2)}) = f(v_2^{(2)}) = \dots = f(v_k^{(2)})$ and $f(v_1^{(1)}) = f(v_1^{(2)})$. Thus, there exists $2 \leq i \leq k$ such that $f(v_1^{(i)}) = 1$. Thus, there are either 0 or $k-1$ vectors among $v_1^{(p)}; v_2^{(p)}; \dots; v_k^{(p)}$ that are colored 1 according to f , contradicting the fact that f respects the PCS $(P; Q)$. \square

Finally, we show the absence of sphere coloring for Lemma 96. Fix integers $k \geq 3; b \geq 0; 1 \leq i \leq k; n \geq 1; k \geq n$. There exists integer n_0 such that for every $n \geq n_0$, there does not exist coloring $f: S^n \rightarrow \{-1, 1\}$ that is folded i.e. $f(-x) = f(x)$, and respects the PCS $(P; Q); P = \text{Ham}_k \{1; k\}; Q = \text{Ham}_k \{0; 1; \dots; k-n\}$.

We dedicate the rest of the section to proving Lemma 96. We pick the conjugation of vectors along the same lines as in Lemma 93. Fix $x \in S^n$. The P-conjugation that we study is a set of vectors $v_1; v_2; \dots; v_k$ that is obtained by first sampling $2 \leq i \leq k$ such that

$$i = \begin{cases} 1 & \text{with probability } \frac{k}{2k-2} \\ k & \text{with probability } \frac{k-2}{2k-2} \end{cases}$$

Then, we sample a uniform point from $\text{Ham}_k(i)$. We get the following properties:

1. (First moments) $v_i \cdot v_0 = \frac{k}{2k-2} \frac{2-k}{k} + \frac{k-2}{2k-2} \cdot 1 = 0$ for every $i \in [k]$.
2. (Second moments). For every $i, j \in [k]$, we get

$$v_i \cdot v_j = \frac{k}{2k-2} \frac{k-1}{k} + \frac{k-2}{2k-2} \cdot 1 = \frac{k-3}{k-1}$$

For ease of notation, let $\alpha = \frac{k-3}{k-1}$. Furthermore, by restricting ourselves to vectors v_i that are orthogonal to v_0 , we just focus on P-conjugations that are a set of unit vectors all of whose pairwise inner product is equal to α . We refer to these set of vectors i.e., a set of k unit vectors $v_1; v_2; \dots; v_k \in S^{n-1}$ an α -conjugation if the inner product of every pair of them is equal to α . Given the folded sphere coloring f our goal is to show that there is a conjugation of vectors V among which exactly $\frac{k-1}{2}$ of them are assigned 1.

Unlike the earlier studied PCSPs, here, the setting when $\alpha = 0$ is relatively straightforward, simply because $\alpha = 0$ implies that there are an arbitrarily large number of unit vectors (as we can pick n to be large enough) all of whose pairwise inner product is equal to α . In particular, we pick a set of $2k-1$ unit vectors all of whose pairwise inner product is equal to α . Among those, k of them are colored the same according to f . By taking the negation of these if needed, we can infer that there are conjugations that are all colored 1, and also α -conjugations that are all colored -1.

Before delving further, we handle the case when $\alpha = 0$ i.e., when $k = 3$. In this case, we just pick a set of k unit vectors that are all orthogonal to each other and their negations. Note that these are $2k$ pairwise orthogonal vectors where exactly $\frac{k-1}{2}$ of them are colored 1 according to f .

Thus, we can pick pairwise orthogonal vectors from this set where exactly k of them are colored +1 according to \mathbf{c} . Henceforth, we assume that $\mathbf{c} \cdot \mathbf{0}$.

To show that there are configurations that have exactly k vectors that are colored +1, we show a connectivity lemma (Lemma 99) where we prove that between any two configurations, there exists a path using only $(k+1)$ -configurations where we change a single vector at each step in the path. As there is an $(k+1)$ -configuration where all $k+1$ vectors are colored +1, and the $(k+1)$ -configuration obtained by negating these vectors where all the vectors are colored -1, the connectivity lemma then shows that for every $\mathbf{c} \in \{-1, 0, 1\}^k$, there exists an $(k+1)$ -configuration that has exactly k vectors that are colored +1.

We first prove the following simplified version of the connectivity lemma that we use to prove Lemma 99.

Lemma 97. Given an $(k+1)$ -configuration $U = (u_1; u_2; \dots; u_{k+1}) \in S^n$, and a unit vector $w \in S^n$ that is orthogonal to each vector u_i , there exists $L := L(k; \mathbf{c})$ and a set of $(k+1)$ -configurations $V_1; V_2; \dots; V_L$ such that

1. The consecutive configurations differ in a single vector v_i , $V_{i+1} \cdot v_i = k - 1$ for every $i \in [L - 1]$.
2. Final configuration contains w i.e., $w \in V_L$, and the initial configuration V_1 is equal to U .

Proof. We prove the lemma by studying the inner product of V with an $(k+1)$ -configuration V , which is equal to all zeroes initially when $V = U$, and changing V one vector at a time such that the inner product of V with w eventually reaches all. Towards this end, for an $(k+1)$ -configuration V , we define the matrix $I(V; w) = (k+1) \times (k+1)$ matrix $I(V; w) = [v_1 v_2 \dots v_k w]^T [v_1 v_2 \dots v_k w]$ defined as follows:

$$I(V; w)_{ij} = \begin{cases} v_i \cdot v_j & \text{if } 1 \leq i, j \leq k \\ v_i \cdot w & \text{if } i = k+1; 1 \leq j \leq k \\ w \cdot v_j & \text{if } j = k+1; 1 \leq i \leq k \\ w \cdot w & \text{if } i = j = k+1 \end{cases}$$

Starting with $I(V; w)$ where $V = U$, our goal is to change one vector v_i at a time so that we eventually reach a configuration where the last column of $I(V; w)$ is equal to $(w \cdot v_1; \dots; w \cdot v_k; 1)$. Note that changing one vector v_i corresponds to changing a single value in the last column (and the corresponding value in the last row) of $I(V; w)$. We show that the opposite direction also holds i.e., by changing a single value in the last column (and the corresponding value in the last row) of $I(V; w)$, we obtain a new matrix that is $I(V^0; w)$ with V^0 being different from V only in a single vector, as long as the new matrix is positive semidefinite.

Claim 98. Suppose that A is a $m \times m$ real symmetric positive semidefinite matrix with $A = U^T U$ with $U = [u_1 u_2 \dots u_m]$ where $u_i \in \mathbb{R}^n$ with $n \geq m$, and A^0 is another real symmetric positive semidefinite matrix such that A^0 and A differ only in $A_{1,2} = A_{2,1}$. Then, there exists $U^0 = [u_1^0 u_2 \dots u_m]$ such that $A^0 = (U^0)^T U^0$.

Proof. As A^0 is a positive semidefinite matrix, there exist $v_1; v_2; \dots; v_m \in \mathbb{R}^n$ such that $A^0 = V^T V$ where $V = [v_1 v_2 \dots v_m]$. Let $A[2:m]$ is a $(m-1) \times (m-1)$ submatrix of A excluding

the first row and the column i.e. $A[2:m]_{ij} = A_{i+1;j+1}$. We define the corresponding submatrix of A^0 as $A^0[2:m]$. Note that $A[2:m] = A^0[2:m]$. Let $U[2:m] = [u_2 u_3 \dots u_m]$, and similarly, let $V[2:m] = [v_2 v_3 \dots v_m]$. Note that $A[2:m] = U[2:m]^T U[2:m]$, and $A^0[2:m] = V[2:m]^T V[2:m]$. However, as $A[2:m] = A^0[2:m]$, we can infer that there exists a unitary $n \times n$ matrix H such that $U[2:m] = HV[2:m]$. Now, setting $U^0 = HV$, we get the required matrix U^0 with $A^0 = (U^0)^T U^0$. \square

Thus, our goal is to obtain a series of $(k+1) \times (k+1)$ real symmetric positive semidefinite matrices $M_1; M_2; \dots; M_L$ such that

1. $M_1 = I(U; w)$.
2. The diagonal entries of M_i for every $i \in [L]$ are all equal to d .
3. All the off diagonal entries of M_L are equal to $\frac{1}{k}$.
4. For every $i \in [L-1]$, M_i and M_{i+1} differ only in one element in the last column (and the corresponding element in the last row).

Towards this end, for $0 \leq k \leq d$, and $d \in [k]$, we define the $(k+1) \times (k+1)$ matrix $M(k; ; d)$ as follows:

$$M(k; ; d)_{ij} = \begin{cases} 1 & \text{if } i = j \\ \frac{1}{k} & \text{if } 1 \leq i, j \leq k \\ \frac{1}{k} + \frac{d}{k} & \text{if } i = k+1; 1 \leq j \leq k \text{ or } j = k+1; 1 \leq i \leq k \\ \frac{1}{k} + \frac{d}{k} & \text{if } i = k+1; d+1 \leq j \leq k \text{ or } j = k+1; d+1 \leq i \leq k \end{cases}$$

Note that $I_{U;w} = M(0; 0; k)$, and our goal is equal to $M(k; 0; k)$. We define the sequence of positive semidefinite matrices $M(0; 0; k); M(0; ; 1); M(0; ; 2); \dots; M(0; ; k); M(k; ; 1); M(k; ; 2); \dots; M(k; ; k)$. Note that at each step, we change a single element in the last column (and the corresponding element in the last row).

The final step is to show that when we set $\frac{1}{k}$, $M(k; ; d)$ is positive semidefinite for every $d \in [k]; 0 \leq k \leq d$. This follows from a simple calculation.

$$\begin{aligned} x^T M(k; ; d)x &= \sum_{i=1}^{k+1} (x_i)^2 + \sum_{i=1}^k \left(\frac{1}{k} + \frac{d}{k} \right) x_i^2 + \sum_{i=1}^k \left(\frac{1}{k} + \frac{d}{k} \right) x_{k+1}^2 + \sum_{i=1}^k \left(\frac{1}{k} + \frac{d}{k} \right) x_i^2 \\ &= \left(\frac{1}{k} + \frac{d}{k} \right) \sum_{i=1}^k x_i^2 + \left(\frac{1}{k} + \frac{d}{k} \right) x_{k+1}^2 + \sum_{i=1}^k \left(\frac{1}{k} + \frac{d}{k} \right) x_i^2 \\ &= \frac{1}{k} \sum_{i=1}^k x_i^2 + \frac{d}{k} \sum_{i=1}^k x_i^2 + \left(\frac{1}{k} + \frac{d}{k} \right) x_{k+1}^2 + \sum_{i=1}^k \left(\frac{1}{k} + \frac{d}{k} \right) x_i^2 \geq 0 \end{aligned}$$

Now, we prove the connectivity lemma.
Lemma 99. Fix an integer $k \geq 2$ and $0 < \epsilon < 1$. Suppose that U and V are two ϵ -conjugations in S^n . Then, there exists $n_0 := n_0(k; \epsilon)$, and $L := L(k; \epsilon)$ such that as long as $n \geq n_0$, there exist ϵ -conjugations $V_1; V_2; \dots; V_L$ such that

1. The end points are U and V i.e., $U = \{v_1\}; V = \{v_L\}$.
2. Any two consecutive configurations differ in exactly one vector v_j , $v_{i+1} = v_i \oplus v_j$ for every $i \in [L-1]$.

Proof. We use induction on k . First, we consider the case when $k=2$. Let $U = \{u_1, u_2\}$, and $V = \{v_1, v_2\}$ be two configurations. Consider an arbitrary vector w that is orthogonal to all the vectors in U and V . Such a w is guaranteed to exist when n is large enough. Now, using Lemma 97, we can infer that there exists a configuration $W = \{w, w^0\}$ such that there is a path of length $O(1)$ from U to W , and from V to W . Thus, there exists a path of length $O(1)$ from U to V .

Assume that the proof holds for $k-1$, and we are given the configurations U and V consisting of k vectors each. We choose a vector w that is orthogonal to each of the vectors in U and V . Using Lemma 97, there are configurations $X = \{w, x_1, x_2, \dots, x_{k-1}\}$ and $Y = \{w, y_1, y_2, \dots, y_{k-1}\}$ such that there is a path of length $O_{k-1}(1)$ from U to X and from V to Y . Now, our goal is to show that there is a path from X to Y of length $O_{k-1}(1)$. We achieve this by restricting ourselves to components orthogonal to w of the $(k-1)$ -sized configurations $X^0 = \{x_1^0, x_2^0, \dots, x_{k-1}^0\}$ and $Y^0 = \{y_1^0, y_2^0, \dots, y_{k-1}^0\}$, where $x_i^0 = x_i \oplus w$ for each $i \in [k-1]$ (and similarly for y_i^0). Note that X^0 and Y^0 are $\binom{k-1}{2}$ configurations, and by the induction hypothesis, there exists a path from X^0 to Y^0 using only vectors orthogonal to w . Adding the component along w , we get a path from X to Y of length $O_{k-1}(1)$, finishing the proof. \square

We are now ready to prove Lemma 96.

Proof. Suppose for contradiction that there exists a coloring $S^n = \{f_0, f_1\}$ that is folded and respects the PCSR. Consider an arbitrary set of vectors $v_1, v_2, \dots, v_{2k-1}$ that are all orthogonal to v_0 , and have pairwise inner product $\frac{1}{2}$. Such a set is guaranteed to exist as follows. Let $U = \{v_1, v_2, \dots, v_{k-1}\}$ be a set of $k-1$ vectors among these that are all assigned the same color. Let these form the configuration U , and the set of negations of these vectors be the configuration V . Using Lemma 99, there exists a path from U to V where we change a single vector in each step. Note that the endpoints of the path have k vectors assigned ± 1 respectively. Since we change at most one vector at a time, there exists a configuration where we have exactly k vectors assigned ± 1 , a contradiction. \square

Lemma 92, Lemma 93, Lemma 94, Lemma 95 and Lemma 96 together with Lemma 84, Theorem 85 and Lemma 86 finish the proof of Theorem 64.

Explicit Construction. We give an explicit construction of an integrality gap instance for $\text{Gap}_{\text{Int}} = (P; Q); P = \text{Ham}_k \{1; k\}; Q = \text{Ham}_k \{0; 1; \dots; k\} \cap \{b\}$ for arbitrary $b \in \{0; 1; \dots; k\}$. Let L be a large constant (depends on k to be set later). We have $n = 2k - 1 + \binom{2k-1}{k} L$ variables $x_i : i \in [2k-1]; x_S^{(i)} : i \in [L]; S \subseteq [2k-1]; |S| = k$. Our constraints are the following: for every subset $S \subseteq [2k-1]$ with $|S| = k$ and $S = \{i_1, i_2, \dots, i_k\}$, we pick L new variables

$x_S^{(1)}; x_S^{(2)}; \dots; x_S^{(L)}$. The constraints are

$$f(x_{i_1}; x_{i_2}; \dots; x_{i_k}); f(x_{i_2}; x_{i_3}; \dots; x_{i_k}; x_S^{(1)}); f(x_{i_3}; \dots; x_S^{(1)}; x_S^{(2)}); \dots; f(x_S^{(L-k+1)}; \dots; x_S^{(L)}); f(x_S^{(L-k+2)}; \dots; x_S^{(L)}; \overline{x_{i_1}}); f(x_S^{(L-k+3)}; \dots; \overline{x_{i_1}}; \overline{x_{i_2}}); \dots; f(\overline{x_{i_1}}; \overline{x_{i_2}}; \dots; \overline{x_{i_k}})$$

We choose ϵ to be the constant factor from Lemma 99 with $\epsilon = \frac{k-3}{k-1}$. The idea is that when all the variables in the constraints $f(x_{i_1}; x_{i_2}; \dots; x_{i_k})$ are all set to be True, all the variables in $f(\overline{x_{i_1}}; \overline{x_{i_2}}; \dots; \overline{x_{i_k}})$, and as there are a series of constraints between them where we alter a single variable, there must exist a constraint where there are exactly $k-1$ variables that are set to True.

Formally, we show that this instance does not satisfy $\text{Ham}_k(f; 0; 1; \dots; k; n)$. Suppose for contradiction that there is an assignment that satisfies all the constraints. Since there are $2k-1$ variables $x_1; x_2; \dots; x_{2k-1}$, at least k of them are set to be true. If not, then at least the negated variables are set to be true. This implies that there is a sequence of constraints where the endpoints are assigned all True and all False, and at every point, we change a single variable. This implies that there is a constraint where there are exactly $k-1$ variables that are set to True, a contradiction.

We now show that the instance has a basic SDP solution with zero error. Set $v_1; \dots; v_{2k-1} \in \mathbb{R}^n$ to the variables $x_1; x_2; \dots; x_{2k-1}$ such that $v_i \cdot v_0 = 0$, and $v_i \cdot v_j = \epsilon$ for every $i \neq j$. Such a set of vectors is guaranteed to exist if n is large enough, and $\epsilon > 0$. Finally, we use Lemma 99 to set the vectors $v_S^{(i)}$ for every $S \subseteq [2k-1]; |S| = k; i \in [L]$.

6.6 The SDP minion

As previously mentioned, polymorphisms are a powerful tool for understanding the computational complexity of PCSPs. However, beyond some of the simplest classes of PCSPs, individually classifying the complexity based on specific polymorphisms can be unwieldy. Instead, one often looks to higher-level structure between classes of polymorphisms, which is captured by the notion of minions (also called clonoids) [Bar+21].

Definition 100. A minion M is a family of sets $M_1; M_2; \dots$, where M_i are the objects of arity i . For every pair of positive integers i and j and map $\pi: [i] \rightarrow [j]$, there exists a map $\pi^*: M_i \rightarrow M_j$ known as a minor map. Further, these minor maps commute ($\pi^* \circ \sigma^* = \sigma^* \circ \pi^*$).

The most commonly discussed minion is the polymorphism minion $\text{Pol}(f)$ of a PCSP (C, f) . In this case the minors correspond to the identification of coordinates. Given a function f of arity i and a map $\pi: [i] \rightarrow [j]$, we have that

$$f^\pi(x_1; \dots; x_j) = f(x_{(1)}; x_{(2)}; \dots; x_{(i)}):$$

It is straightforward to verify that the minor maps commute in the specified manner.

However, the 'objects' of a minion need not correspond to mathematical functions. The following are a few known examples:

- The trivial minion M_{triv} has that every arity has a single object $m_i = f$. All minor maps are thus the "trivial" map.

- The dictator minion (or projection minion) M_{dict} has each $M_i = [i]$. The minor maps are then application of $\pi_k = \pi_{(k)}$.
- The Basic LP minion M_{BLP} has each $M_i = f(p_1; \dots; p_k) : p_1, p_2; \dots; p_k \in [0, 1]; p_1 + \dots + p_k = 1$ be the probability distributions on elements. The minor maps combine elements of the probability distribution which map to the same value. That is,

$$(p_1; \dots; p_k)^\pi = \sum_{(k)=1}^0 \prod_{(k)=j}^1 p_k A :$$

In order to better understand the complexity of PCSPs, we relate the polymorphic minions to other minions like the ones mentioned above. We can determine the relationship between minions via minion homomorphisms

Definition 101. A minion homomorphism $\pi : M \rightarrow N$ between two minions consists of maps $\pi_i : M_i \rightarrow N_i$ such that these maps commute with the respective minor maps of M and N . That is, for all $f \in M_i$ and $\pi : [i] \rightarrow [j]$, we have that

$$\pi_j(f^\pi) = \pi_i(f)$$

The key theorem about the Basic LP minion is the following.

Theorem 102 ([Bar+21]). The Basic LP solves a PCSP f and only if $M_{\text{BLP}} \leq \text{Pol}(f)$.

6.6.1 SDP Minion Definition

Like how the the BLP minion corresponds to probability distributions, the SDP minion we construct corresponds to SDP vectors. Similar techniques have been used in other minion constructions [CZ22a].

Let \mathbb{R}^k be finite sequences of real numbers which are eventually 0 and thus can be thought of as the union $\mathbb{R}^1 \cup \mathbb{R}^2 \cup \mathbb{R}^3 \cup \dots$. It's not hard to see that \mathbb{R}^k is an inner product space.

We now define the minion M_{SDP} whose k -arity symbol is a list of k vectors $(w_1; \dots; w_k)$ in \mathbb{R}^k . When convenient, we shall think of the whole object as a matrix $W \in \mathbb{R}^{k \times k}$.

We impose the following conditions on the vectors:

1. For all $i, j \in [k]$ with $i \neq j$, $w_i \cdot w_j = 0$.
2. $\sum_{i \in [k]} \|w_i\|_2^2 = 1$.

Observe that the second condition is equivalent to $\sum_{i \in [k]} w_i \cdot w_i = 1$.

The minors of M_{SDP} are not too surprising, given $W \in M_{\text{SDP}}^{(k)}$ and a map $\pi : [k] \rightarrow [j]$, W^π is the matrix in $\mathbb{R}^{j \times j}$ where $w_i^\pi := \sum_{j \in \pi^{-1}(i)} w_j$.

Claim 103. M_{SDP} is a minion.

Proof. First, for each $W \in M_{\text{SDP}}^{(k)}$ and a map $\pi : [k] \rightarrow [j]$ we verify that $W^\pi := W^\pi \in M_{\text{SDP}}^{(j)}$.

First, $x_i \in \{0, 1\}$. We have that

$$\begin{aligned} w_i^0 w_{i^0}^0 &= \sum_{j \in [k]} \sum_{i^0 \in [k]} w_j^0 w_{i^0}^0 \\ &= \sum_{j \in [k]} \sum_{i^0 \in [k]} w_j^0 w_{i^0}^0 \\ &= 0 \end{aligned}$$

Further,

$$\sum_{i \in [k]} w_i^0 = \sum_{i^0 \in [k]} w_{i^0}^0 = 1$$

Thus, $W \in \text{SDP}$:

The only remaining condition to check is that minors commute. Consider $[a] \cup [b]$ and $[b] \cup [c]$. Also consider $U \in \text{SDP}^{(a)}$, $V \in \text{SDP}^{(b)}$, and $W \in \text{SDP}^{(c)}$ such that $U = V$ and $W = V$. We seek to verify that $U = W$. For all $i \in [c]$, we have that

$$\begin{aligned} w_i &= \sum_{i^0 \in [k]} v_{i^0} \\ &= \sum_{i^0 \in [k]} u_{i^0} \\ &= \sum_{i^0 \in [k]} w_{i^0} \end{aligned}$$

as desired. □

The goal of this section is to prove the following theorem
 Theorem 104. If the exact basic SDP solves PCSP(A; B) if and only if $M_{\text{SDP}} \approx \text{Pol}(A; B)$

6.6.2 An alternative Basic SDP

We now present a modified basic SDP for which it is easier to make our arguments. We shall use notation similar to that of the [Z22a].

Let $(A; B)$ be a PCSP template. Let \mathcal{X} be an instance of PCSP(A; B). We let \mathcal{A} be the domain of A , etc., so X is the variables of \mathcal{X} . We let R^A be the set of relations of \mathcal{X} (or "clause types"). We let R^X be the constraints of \mathcal{X} .

¹³By 'exact' we mean an SDP solution that perfectly satisfies all the constraints. This is not feasible in principle as infinite precision is needed in the computed SDP matrix (c.f., [Fre04]). In practice, SDP solutions can be computed to poly(n) bits of precision, which we would consider to be a $1=2^{\text{poly}(n)}$ -robust solution. For PCSPs, it is unknown if there is a difference between exact, $1=2^{\text{poly}(n)}$ -robust and $O(1)$ robust.

A basic SDP solution to PCSP(A; B) is a collection of vectors $u_{x;a} \in \mathbb{R}^k$ for all $x \in X; a \in A$ as well as $v_{x;a} \in \mathbb{R}^k$ for $x \in X$ and $a \in s^R(x)$ (which is defined to be the set of the valid assignments to the clause) with the following properties:¹⁴

1. For all $x \in X$, $U_x := (u_{x;a} : a \in A) \in M_{SDP}^{(A)}$:
2. For all $x \in X$, $V_x := (v_{x;a} : a \in s^R(x)) \in M_{SDP}^{(A)}$.
3. For all $x \in X$ with arity k and $i \in [k]$ and $a \in A$, we have that

$$u_{x_i;a} = \sum_{a \in s^R(x)} v_{x;a} \cdot \mathbb{1}_{a_i = a}$$

Equivalence with traditional basic SDP

The traditional basic SDP does not specify vectors. Rather, the traditional basic SDP keeps track of a probability distribution μ on assignments to \mathcal{X} . And for any pairs of variable, assignment pair $(x_i; a_i); (x_j; a_j)$, we have that

$$u_{x_i;a_i} \cdot u_{x_j;a_j} = \Pr_a [a_i = a_i \wedge a_j = a_j]:$$

Clearly any solution to our Basic SDP is a solution to that basic SDP, as the probability is precisely $\sum_{a \in s^R(x)} \mu(a)$, and the above linear condition checks out as a property of the vectors.

The other direction is more tricky. Assume we have a traditional basic SDP solution. Recall that each $u_{x;a} \in \mathbb{Q}^k$. Since there are finitely many vectors, there exists some $N \in \mathbb{N}$ such that each $u_{x;a}$ has its support within the first N coordinates. Provisionally assign $u_{x;a}$ to be $\sum_{i=1}^N \mu_{x;a}(a_i) e_i$, where $e_i > 0$ are chosen uniquely for each clause-assignment pair. As there are finitely many clause-assignment pairs, all of these are in \mathbb{Q}^k .

However, it is obvious the u 's won't be compatible with the v 's as they are in trivially-intersecting vector spaces. However, we can define over all suitable choices $\phi \in \mathbb{Q}^k$ and $a \in A$,

$$\hat{u}_{x_i;a_i} := \sum_{a \in s^R(x)} \hat{v}_{x;a} \cdot \mathbb{1}_{a_i = a}$$

It is not hard to see that $u_{x_i;a_i} \cdot u_{x_j;a_j} = \hat{u}_{x_i;a_i} \cdot \hat{u}_{x_j;a_j}$ for all suitable choices of $x_i; a_i; x_j; a_j$.

For a fixed x , let $\hat{V}_x \subseteq \mathbb{Q}^k$ be the subspace spanned by $\hat{u}_{x;a}$ and $\hat{v}_{x;a}$ for the $x \in X$. Note that each \hat{u} is also in \hat{V}_x . Since the \hat{u} 's and \hat{v} 's have the same dot products, there is a rigid rotation $\rho : \hat{V}_x \rightarrow \hat{V}_x$ which sends each $\hat{u}_{(x;a)}$ to $u_{(x;a)}$.

Define $v_{x;a} = \rho(\hat{v}_{x;a})$. Then, observe that

$$u_{x_i;a_i} = \rho(\hat{u}_{x_i;a_i}) = \sum_{a \in s^R(x)} \rho(\hat{v}_{x;a}) \cdot \mathbb{1}_{a_i = a} = \sum_{a \in s^R(x)} v_{x;a} \cdot \mathbb{1}_{a_i = a}$$

as desired. Thus, the two SDP formulations are equivalent.

¹⁴Usually, there is a special vector u_0 , but we can omit it without changing the power of the algorithm. This will be more convenient for the analysis.

6.6.3 From minion homomorphism to SDP rounding algorithm

First we show the "easy" direction that the minion homomorphism implies that the basic SDP solves PCSP(A; B).

Theorem 105. If $M_{SDP} \models \text{Pol}(A; B)$, then the basic SDP solves PCSP(A; B).

Proof. Fix X and an exact SDP solution M^X with corresponding vectors $u_{x;a}$ and $v_{x;a}$ (and $U_x; V_x$) with the prescribed properties. Let: $M_{SDP} \models \text{Pol}(A; B)$ be the minion homomorphism.

Define the assignment $f: X \rightarrow B$ to be $(U_x)(a: a \in A)$. Unpacking this, $(U_x) \in \text{Pol}(A; B)$ is of arity A , so we just plug in the coordinates a listed in a canonical order. For any clause c , we know that $b := (V_x)(a: a \in s^R(x))$ satisfies R^B . Thus, it suffices to show that $b_i = f(x_i)$ for all $i \in [k]$.

Let $\pi_i: \mathbb{R}^A \rightarrow \mathbb{R}^A$ which maps a to a_i . It is straightforward to show that condition (3) of the basic SDP implies that $u_x = (V_x)_{\pi_i}$. Thus, since π_i is a minion homomorphism,

$$f(x_i) = (U_x)(a: a \in A) = (V_x)_{\pi_i}(a: a \in A) = (V_x)_i(a: a \in s^R(x)) = b_i;$$

as desired. Thus, the exact basic SDP solves PCSP(A; B). \square

6.6.4 From SDP rounding algorithm to minion homomorphism

We now prove the converse.

Theorem 106. If the basic SDP solves PCSP(A; B), then $M_{SDP} \models \text{Pol}(A; B)$

Proof. We adopt the proof technique of [Z22a]. Let $F \subseteq M_{SDP}$ be any finite subset. Let F be an instance of PCSP(A; B) whose variables $F = F \setminus M_{SDP}^{(A)}$ – the arity A elements of F . The clauses R^F are k -tuples $(W_1; \dots; W_k)$ of F for which there is $Z \in F \setminus M_{SDP}^{(R^A)}$ with the following property. For all $i \in [k]$, let $\pi_i: \mathbb{R}^A \rightarrow \mathbb{R}^A$ be the i th coordinate projection map. Then, $W_i = Z_{\pi_i}$.

If we can show that the rational basic SDP solves F , then we know that $F \models B$ (by definition of the rational basic SDP solving PCSP(A; B)). Then, via an argument (e.g., like the De Bruijn-Erdős theorem [BE51], for more details see Remark 7.13 of [Bar+21] [Z22a], etc.) this implies that free structure $F_{M_{SDP}}^{(A)} \models B$ which implies that $M_{SDP} \models \text{Pol}(A; B)$. Thus, it suffices to construct a rational basic SDP solution F to

The remainder of the proof writes itself. For every $W \in F$ and $a \in A$, let $u_{W;a} = w_a$ (the a th column of W). Likewise, for every clause $(W_1; \dots; W_k)$ via $Z \in F \setminus M_{SDP}^{(R^A)}$ and $a \in s^R(Z)$ (that is a valid solution to the clause indexed by Z), we let $v_{W;a} = Z_a$.

We then need to check conditions 1-3. Conditions 1 and 2 are verbatim from $W; Z \in M_{SDP}$. Condition 3 is precisely that $W_i = Z_{\pi_i}$.

Since all the vectors have rational coordinates, the dot product matrix of these vectors is a rational basic SDP solution. That is, the rational basic SDP solves F . \square

¹⁵See [Bar+21] for a precise definition.

This completes the proof of Theorem 104

6.7 Missing Proofs

Proof of Lemma 79. We first consider the case when $n = 1$. Without loss of generality, let $P = Q = \{f+1, g\}$, and we use $v_1 = v_0 + v_1^0$ to denote the SDP vector corresponding to the variable used in the constraint. As the basic SDP has error at most $\frac{1}{r}$, we get that

$$\|v_1 - v_1^0\| \leq \frac{1}{r}$$

As $\|v_1\|_2^2 + \|v_1^0\|_2^2 = 1$, $\|v_1^0\|_2 = O(r^{-0.25})$. Thus, using Proposition 71, we get that $\|v_1^0\|_2 = O(r^{-0.25})$ with probability at least $1 - e^{-\frac{r^2}{2}}$. On the other hand, using Proposition 70, we get that $\|v_1 - v_1^0\|_2 \leq \frac{1}{r}$ with probability at least $1 - \frac{1}{r}$. This implies that

$$\|v_1 - v_1^0\|_2 \leq \frac{1}{2r}$$

Thus, with probability at least $O(\frac{1}{r})$, we have

$$\|v_1 - v_1^0\|_2 = O(r^{-0.25}) < \frac{1}{2r} = \|v_1 - v_1^0\|_2$$

Hence, with probability at least $O(\frac{1}{r})$, we round the variable to 1.

We now consider the general case when $n \geq 2$. Note that the above proof for $n = 1$ holds when $P = \text{Ham}_k\{0, g\}$ or when $P = \text{Ham}_k\{f, k, g\}$. We are left with the setting when $P = \text{Ham}_k\{0, k, g\}$. In order to show that our algorithm is a robust algorithm for this PCSP, it suffices to show that all the elements in the predicates are rounded to the same value with high probability. Consider k . We show that the probability that the variables x_i and x_j get rounded to different values is at most $O(\frac{1}{r})$. Using the union bound over all the $\binom{k}{2}$ pairs of indices, we get our required claim.

We first collect useful properties using the fact that the basic SDP is supported with probability at least $1 - c$ on P .

1. (First moment.) We have

$$\|v_i - v_j\|_2 \leq 2c$$

2. (Second moment.) We have

$$\|v_i - v_j\|_2 \leq 1 + 2c$$

Using this, we get

$$\begin{aligned} \|v_i^0 - v_j^0\|_2^2 &= \|v_i - v_j + (v_i - v_i^0) - (v_j - v_j^0)\|_2^2 \\ &\leq \|v_i - v_j\|_2^2 + \|v_i - v_i^0\|_2^2 + \|v_j - v_j^0\|_2^2 \\ &\leq (1 + 2c)^2 + O(c) + O(c) \\ &= O(c) \end{aligned}$$

Thus, $\|v_i^0 - v_j^0\|_2 = O(\sqrt{c})$.

As earlier, we assume that at most P^{-} .

Recall that our goal is to upper bound the probability that x_i and x_j are rounded to different values. Without loss of generality, suppose that x_i is rounded to -1 , and x_j is rounded to 1 . We get that

$$\begin{aligned} h; v_i^0 & \leq jh; v_{0ij} \\ h; v_j^0 < & jh; v_{0ij} \end{aligned}$$

Using Proposition 71, we can infer that

$$jh; v_i^0 \leq h; v_j^0 \leq O(r^{P-\bar{c}})$$

with probability at least $\frac{1}{r}$.

We consider two cases: first, when $n_{ij} \leq \frac{1}{2}$. As $\frac{1}{r} + kv_i^0 k_2^2 = 1$, and $j - i \leq j - 2c$, we get that $v_j^0 = (1)$. In this case, we have

$$\begin{aligned} h; v_j^0 & \leq h; v_i^0 \leq O(r^{P-\bar{c}}) \\ jh; v_{0ij} & \leq O(r^{P-\bar{c}}) \end{aligned}$$

Thus, we have

$$h; v_j^0 \leq [jh; v_{0ij} \leq O(r^{P-\bar{c}}); jh; v_{0ij}]$$

Here, $h; v_j^0 \leq [p; q]$ where $q - p = O(r) + O(r^{P-\bar{c}}) = O(r)$. However, as $v_j^0 = (1)$, this happens with probability at most $O(r)$.

Now, suppose that $n_{ij} > \frac{1}{2}$. We have

$$jh; v_{0ij} \leq jh; v_{0ij} \leq 2cjh; v_{0ij}$$

However, as $h; v_{0i} \leq N_p(0; 1)$, we have $jh; v_{0ij} \leq r$ with probability at least P^{-} . Thus, with probability at least P^{-} , we have

$$h; v_i^0 \leq jh; v_{0ij} \leq h; v_j^0 \leq 2cr$$

We have $jh; v_{0ij} \leq [p; q]$ where $q - p = O(cr) + O(r^{P-\bar{c}})$. However, this happens with probability at most $O(\frac{r^{P-\bar{c}}}{r}) = O(\frac{1}{r})$. \square

Proof of Lemma 82 Suppose that $a = \text{sgn}(x - y)$ for $x, y \in A(P)$, and $x_i \in y_i \in \mathbb{Z}[k]$. By modifying the above combinations slightly, we can assume that x and y are a rational combination of P while still preserving the fact that $a = \text{sgn}(x - y)$. In other words, there exist $p_1, p_2, \dots, p_K, q_1, q_2, \dots, q_K \in \mathbb{Z}$ such that $\sum_{i \in [K]} p_i = \sum_{i \in [K]} q_i = 1$, and $x = \sum_{i \in [K]} p_i a_i$, $y = \sum_{i \in [K]} q_i a_i$, where $\sum_{i \in [K]} a_i = 1$. Let N be a positive integer such that we can write $p_i = \frac{p_i^0}{N}$, $q_i = \frac{q_i^0}{N}$ where p_i^0, q_i^0 are integers for every $i \in [K]$.

Let S be a multiset of $\sum_{i \in [K]} a_i$ where we take union over all $i \in [K]$, p_i^0 copies of a_i , assign them a sign $\text{sgn}(p_i^0)$, and q_i^0 copies of a_i , assign them a sign $\text{sgn}(q_i^0)$. As we have

Let $P = \{x_1, x_2, \dots, x_L\}$ and $Q = \{y_1, y_2, \dots, y_N\}$. We get that there are equal number of vectors in S that are assigned sign $+1$ and equal number of them that are assigned sign -1 . Let z denote the signed sum of all vectors (including repetitions) in S . Note that $\text{sgn}(z) = \text{sgn}(x - y)$. As each element of S is an integer, we get that the absolute value of each coordinate in z is at least 1. Furthermore, we can take multiple copies of S to ensure that the absolute value of each coordinate in z is at least 2. Now, we add an arbitrary element w with sign $+1$ to S . Note that we still have that the signed sum of S is z . i.e., the updated S satisfies $\text{sgn}(z) = \text{sgn}(x - y)$. Furthermore $z = x_1 - x_2 + \dots + x_L$ where each $x_i \in P$. Thus, $a = \text{sgn}(x - y) = \text{sgn}(z) = \text{sgn}(x_1 - x_2 + \dots + x_L) \in O_{AT}(P)$. Thus,

$$f \text{sgn}(x - y) : x, y \in A(P); \exists x_i \in y_i, g \in O_{AT}(P)$$

To prove the other direction, suppose $a \in O_{AT}(P)$. That is, $a = \text{sgn}(x_1 - x_2 + \dots + x_L)$. Let S be a multiset of x_1, x_2, \dots, x_L with the corresponding sign as in the summation. As S is non-trivial in every coordinate i.e., for every $i \in [k]$, there exist assignments $x_i \in P$ where $x_i = +1$, and similarly, $x_i \in P$ where $x_i = -1$. By adding vectors with both signs $+1$ and -1 , we can assume that S is non-trivial in every coordinate while still preserving the fact that the sign vector of the signed sum of S is equal to $\text{sgn}(a)$. We modify S while still preserving this property to ensure that the signed sum of vectors in S has absolute value at least 1 in every coordinate.

As there are odd number of vectors in S , the signed sum of the vectors has absolute value at least 1 in every coordinate. Fix a vector $x \in S$. Create two copies of every other vector in S (with the same sign as the original). Note that this operation does not alter the sign vector of the signed sum of the vectors in S . We can repeat this process at most $|S|$ times to ensure that in the final multiset S , the signed sum has absolute value at least 1 in every coordinate. Finally, we add an arbitrary vector with sign -1 to S , to ensure that there are equal number of vectors with $+1$ and -1 sign in S . Overall, we get that there are $x_1, x_2, \dots, x_N \in P$ and $y_1, y_2, \dots, y_N \in P$ such that

$$a = \text{sgn}(x_1 + \dots + x_N - y_1 - \dots - y_N) = \text{sgn} \left(\frac{1}{N}x_1 + \dots + \frac{1}{N}x_N - \frac{1}{N}y_1 - \dots - \frac{1}{N}y_N \right)$$

Thus, we get that $f \text{sgn}(x - y) : x, y \in A(P); \exists x_i \in y_i, g \in O_{AT}(P)$, completing the proof that

$$O_{AT}(P) = f \text{sgn}(x - y) : x, y \in A(P); \exists x_i \in y_i, g \in O_{AT}(P) \quad \square$$

Proof of Lemma 84. We extensively use the properties of MAJ polymorphisms of symmetric PCSPs, and ppp-reductions between symmetric folded PCSPs proved in [BG21b].

Consider $(P; Q) \in \text{PCSP}$ be of arity k such that $\text{MAJ}_{L_1}; \text{AT}_{L_2} \not\leq \text{Pol}(\cdot)$ for some odd integers L_1, L_2 . Note that $\text{MAJ}_{L_1} \leq \text{Pol}(\cdot)$ as in that case $\text{MAJ}_{L_1}(P) = P$. Q , contradicting the fact that $\text{MAJ}_{L_2} \not\leq \text{Pol}(P; Q)$. Thus, there exists $b \in \{1, 2, \dots, k-1\}$ such that $\text{Ham}_k \leq \text{Pol}(P)$.

Case 1. We first consider the case when $P = \text{Ham}_k \leq \text{Pol}(P; Q)$. As $\text{MAJ}_{L_1} \leq \text{Pol}(P; Q)$, there exists $b \in \{0, 1, \dots, k\}$ such that $\text{Ham}_k \leq \text{Pol}(P; Q)$ and $\text{Ham}_k \leq \text{Pol}(P; Q)$.

Suppose that $b \neq 0, k$. Let $Q^0 = \{x_1, \dots, x_k\} \in \text{Ham}_k \leq \text{Pol}(P; Q)$. By definition, $\text{MAJ}_{L_1} \leq \text{Pol}(P; Q^0)$. As $O_{AT}(\text{Ham}_k \leq \text{Pol}(P; Q)) = \text{Ham}_k \leq \text{Pol}(P; Q^0)$, we get that $\text{AT}_{L_2} \leq \text{Pol}(P; Q^0)$. Thus, we get $(P; Q)$

where $P = \text{Ham}_k f l g, Q = \text{Ham}_k f 0; 1; \dots; k g n f b g$ where $b \geq 1; 2; \dots; k - 1 g n f l g$. Note that $\text{MAJ}; \text{AT} \not\geq \text{Pol}(P; Q)$. We now relax this PCSP furthermore, updating $P, Q; l; k; b$ while preserving the following two properties:

1. At every step $P = \text{Ham}_k f l g; Q = \text{Ham}_k f 0; 1; \dots; k g n f b g$ where $b \geq 1; 2; \dots; k - 1 g n f l g$.
2. $\text{MAJ}; \text{AT} \not\geq \text{Pol}(P; Q)$.

As $O_{\text{MAJ}}(P) = \text{Ham}_k f 0; 1; \dots; k g \setminus f 2l - k + 1; \dots; 2l - 1 g$, and $b \geq O_{\text{MAJ}}(P)$, we get that $b \geq 2l - k + 1; \dots; 2l - 1 g \setminus f 0; \dots; k g$. Furthermore, $ab > 0$, we get that $a > 1$. Similarly, we get that $a < k - 1$. This also implies that $4 \leq a \leq 2l - 1; \dots; k - 1 g$.

We use the following two tools to relax the PCSP:

1. Given a PCSP $P = \text{Ham}_k f l g; Q = \text{Ham}_k f 0; 1; \dots; k g n f b g$ where $b \geq 1; 2; \dots; k - 1 g n f l g$, then the PCSP $P^0 = \text{Ham}_k f \lfloor l g; Q = \text{Ham}_k f 0; 1; \dots; k g n f b g$ is a relaxation of $(P; Q)$ (Claim 4:2 of [BG21b]). As long as $a \leq k - 1$ and $b \in 2l - k + 1, \dots, k - 1$, this relaxation preserves the above two properties.
2. Given a PCSP $P = \text{Ham}_k f l g; Q = \text{Ham}_k f 0; 1; \dots; k g n f b g$ where $b \geq 1; 2; \dots; k - 1 g n f l g$, then the PCSP $P^0 = \text{Ham}_k f \lceil l g; Q = \text{Ham}_k f 0; 1; \dots; k g n f b g$ is a relaxation of $(P; Q)$ (Claim 4:4 of [BG21b]). As long as $a > 1$ and $b \in 2l - 1, \dots, k - 1$, this relaxation preserves the above two properties.

Now, we relax the PCSP using the above two steps. As a is decreasing at every step, this procedure terminates at some point. Then, either of the two conditions hold:

1. $b = 1; b = 2l - k + 1$. In this case, we get that $a = \frac{k}{2}$ and $b = 1$. Thus, $(P^0; Q^0)$ is a relaxation of $(P; Q)$ where $P^0 = \text{Ham}_k f \frac{k}{2} g; Q = \text{Ham}_k f 0; 1; \dots; k g n f 1 g$ where k is even and l is at least $\frac{k}{2}$.
2. $b = k - 1; b = 2l - 1$. In this case, we get that $a = \frac{k}{2}$ and $b = k - 1$. Thus, $(P^0; Q^0)$ is a relaxation of $(P; Q)$ where $P^0 = \text{Ham}_k f \frac{k}{2} g; Q = \text{Ham}_k f 0; 1; \dots; k g n f k - 1 g$ where k is even and l is at least $\frac{k}{2}$.

Suppose that there is no $l \geq 0; k g$ such that $\text{Ham}_k f l g \in O_{\text{MAJ}}(P) \cap Q$. As $O_{\text{MAJ}}(P) \cap Q$, by negating the variables if needed, we can assume that $\text{Ham}_k f 0 g \in O_{\text{MAJ}}(P) \cap Q$. Furthermore, there exists $b \geq 1; 2; \dots; k - 1 g$ such that $\text{Ham}_k \cap Q$ as $O_{\text{AT}}(P) \cap Q$. Thus, we obtain a relaxation $(P; Q)$ of the original PCSP such that $P = \text{Ham}_k f l g; Q = \text{Ham}_k f 1; \dots; k g n f b g$ where $l; b \geq 1; 2; \dots; k - 1 g; b > 2l - 1$. By using the first type of relaxation used above (Claim 4:2 of [BG21b]), we obtain a new relaxation such that $P = \text{Ham}_k f l g; Q = \text{Ham}_k f 0; 1; \dots; k g n f 0; k - 1 g$ where $l \geq 1; 2; \dots; k - 1 g; l \leq \frac{k-1}{2}$.

Case 2 There exist $l \geq 0$ such that $\text{Ham}_k f l g \in P$. Recall that $P \cap \text{Ham}_k f 0; k g$. This implies that $O_{\text{AT}}(P) = \text{Ham}_k f 0; 1; \dots; k g$. Hence, we can get a relaxation $(P; Q^0)$ of the original PCSP such that $Q^0 = \text{Ham}_k f 0; 1; \dots; k g n f b g$ such that $\text{Ham}_k f b g \in O_{\text{MAJ}}(P)$.

For ease of notation, let $P = \text{Ham}_k S$, where $S \in \{0; 1; \dots; k\}$. First, consider the case when $\min S = 0; \max S = k$. As mentioned earlier, we know that there exists $b \in \{1; 2; \dots; k - 1\}$ such that $\text{Ham}_k f b g \in P$. Thus, we can reduce the existing PCSP $(P; Q)$ where $P =$

$\text{Ham}_k f 0; l; kg; Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$ where $b \geq f 0; l; kg$. We consider three cases separately:

1. Suppose that $l = \frac{k-1}{2}$. In this case, we have a relaxation $(P; Q)$ where $P = \text{Ham}_k f l; kg$ and $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$ where $b \geq f l; kg$. Note that this relaxation does not contain AT or MAJ as polymorphisms.
2. Suppose that $l = \frac{k}{2}$. In this case, we have a relaxation $(P; Q)$ where $P = \text{Ham}_k f l; kg$ and $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$ where $b \geq f 0; l; kg$. This also doesn't have AT and MAJ as polymorphisms. Furthermore, we have shown earlier that we can relax this further to earlier mentioned three PCSPs.
3. Suppose that $l = \frac{k+1}{2}$. In this case, we have a relaxation $(P; Q)$ where $P = \text{Ham}_k f 0; l; kg$ and $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$ where $b \geq f 0; l; kg$. Note that this relaxation does not contain AT or MAJ as polymorphisms.

Thus, we have a relaxation $(P; Q)$ of the original PCSP where $P = \text{Ham}_k f l; l^0; kg$; $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$ such that $l < l^0$; $f l; l^0; kg \in f 0; kg$, $b \geq O_{\text{MAJ}}(P)$. We end up with the same relaxation when $f \min S; \max Sg \in f 0; kg$.

If $f l; l^0; kg = f 1; kg$ or $f 0; k - 1; kg$, we get a relaxation of the original PCSP where $P = \text{Ham}_k f 1; kg$; $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$, and we are done. If not, we get a series of relaxations of the original PCSP maintaining the two below properties:

1. $P = \text{Ham}_k f l; l^0; kg$ with $l < l^0$ and $f l; l^0; kg \in f 0; kg$ and $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$. We also assume that $f l; l^0; kg \in f 1; kg$ and $f l; l^0; kg \in f 0; k - 1; kg$.
2. $\text{AT}; \text{MAJ} \not\subseteq \text{Pol}(P; Q)$.

As with the earlier case, we update the PCSP using the two relaxations below:

1. We get $P^0 = \text{Ham}_k f l; l^0; kg$ and $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$ using Claim 4:2 of [BG21b]. For this to be a valid relaxation preserving the above properties, we need that $b \in k$ and $b \in 2l - k + 1$.
2. We get $P^0 = \text{Ham}_k f l - 1; l^0 - 1; kg$ and $Q = \text{Ham}_k f 0; 1; ::::; kg n fbg$ using Claim 4:4 of [BG21b]. For this to be a valid relaxation preserving the above properties, we need that $l \in 0$; $b \in 0$ and $b \in 2l^0 - 1$.

As the arity of the predicates is decreasing at each step, this process terminates in finite steps. When we are unable to obtain a new relaxation using the above procedures, one of the following must be true.

1. $l^0 = k$; $b = 0$. In this case, we have a PCSP $(P; Q)$ where $P = \text{Ham}_k f l; kg$, $Q = \text{Ham}_k f 1; 2; ::::; kg$, where $l \in 0$; $l = \frac{k-1}{2}$.
2. $b = k$; $l = 0$. This can be relaxed to the above by negating the variables.
3. $b = k$; $b = 2l^0 - 1$. We have $l^0 = \frac{k+1}{2}$. In this case, we have PCSP $(P; Q)$ where $P = \text{Ham}_k f l; \frac{k+1}{2}; kg$, $Q = \text{Ham}_k f 0; 1; 2; ::::; k - 1; kg$, where $l = \frac{k-1}{2}$.
4. We have $b = 2l - k + 1$; $b = 0$. In this case, we can relax to the above by negating the variables.

Thus, we have relaxed the original PCSP into either of the following PCSPs.

1. k is even, and $P = \text{Ham}_k f \lfloor \frac{k}{2} \rfloor g$; $Q = \text{Ham}_k f 0; 1; \dots; k \lfloor \frac{n}{2} \rfloor b g$ where $2 \leq f \leq k-1$.
2. k is odd, $P = \text{Ham}_k f \lfloor \frac{k+1}{2} \rfloor g$, $Q = \text{Ham}_k f 0; 1; 2; \dots; k-1 g$, where $1 \leq f \leq \frac{k-1}{2}$.
3. $P = \text{Ham}_k f \lfloor \frac{k}{2} \rfloor g$, $Q = \text{Ham}_k f 1; 2; \dots; k g$, where $1 \leq f \leq \frac{k-1}{2}$.
4. $P = \text{Ham}_k f \lfloor \frac{k}{2} \rfloor g$; $Q = \text{Ham}_k f 0; 1; \dots; k \lfloor \frac{n}{2} \rfloor b g$; $k-1 g$ where $2 \leq f \leq k-1$; $1 \leq b \leq \frac{k-1}{2}$.
5. $P = \text{Ham}_k f 1; k g$; $Q = \text{Ham}_k f 0; 1; \dots; k \lfloor \frac{n}{2} \rfloor b g$ for arbitrary b .

□

Chapter 7

Revisiting Alphabet Reduction

7.1 Introduction

Constraint Satisfaction Problem (CSP) is a canonical NP-complete problem. Assuming NP, no polynomial time algorithm can find a satisfying assignment to a satisfiable CSP instance. If we are happy with the easier goal of satisfying a $\alpha(1)$ fraction of constraints, does there exist an efficient algorithm to do so? Answering this in the negative, the fundamental PCP theorem [Aro+98; AS98] implies that for some fixed integers $k, q \geq 2$ and $c < 1$, it is NP-hard to find an assignment satisfying a fraction α of constraints in a satisfiable CSP of arity k over alphabet $\Sigma = \{0, 1, \dots, q-1\}$. Further this result holds for the combinatorial cases $(k, q) = (2, 3)$ and $(3, 2)$. The PCP theorem lies at the center of a rich body of work that has yielded numerous inapproximability results, including many optimal ones.

The PCP theorem was originally proved using algebraic techniques such as the low-degree test and the sum-check protocol. In a striking work, Dinur [Din07] gave an alternative combinatorial proof of the PCP theorem. Her proof works by amplifying the 'Unsatisfied value' of a CSP instance — the fraction of constraints any assignment should violate. The goal is to show that it is NP-hard to distinguish if the Unsatisfied value of a CSP instance is equal to 0 or at least a constant $\alpha > 0$. Starting with a NP-hard problem such as coloring with m constraints, we can already deduce that it is NP-hard to identify if Unsatisfied value is equal to 0 or at least $1/m$. The Unsatisfied value is increased slowly and iteratively via two steps: gap amplification and alphabet reduction. In gap amplification, we incur a constant factor blow up in the size of the instance, and get a constant factor improvement in the Unsatisfied value. However, this step also blows up the alphabet size. To alleviate this, alphabet reduction brings back the alphabet size to an absolute constant while losing a constant factor in the Unsatisfied value (and blows up the instance size by a constant factor). Combining both the steps, we can increase the Unsatisfied value by a constant factor (say 2) incurring a constant factor blow up in the size of the instance. Repeating this $\log(1/\alpha)$ times proves the PCP theorem.

In this chapter, we revisit the alphabet reduction step. Dinur implemented this step by an "inner" PCP construction, which is in effect a gadget reducing a specific predicate to be tested to a collection of constraints over a fixed (say Boolean) alphabet, such that if unsatisfiable,

then a constant fraction of constraints must be violated by any assignment. This inner PCP is then applied to all constraints in the CSP instance G produced by the gap amplification step. The collection of inner PCPs as such only ensure that each constraint is individually satisfiable, which is not very meaningful. To ensure that the inner PCPs together ensure that the constraints of G are all satisfiable by a single consistent assignment, error-correcting codes are used to encode the purported assignments to the variables. The inner PCP is also replaced by an Assignment Tester that ascertains whether the specific assignment given by these encodings satisfies the predicate being checked.

The key observation driving this work is that instead of designing the inner PCP for arbitrary constraints (as in Dinur's paper), we can first reduce the CSP instance G produced by gap amplification to a Label Cover instance. Label Cover is a special kind of CSP which has arity 2, and whose underlying relations are functions (so the value of one of the variables in each constraint is determined by the value taken by the other variable in that constraint). Conveniently for us, we also observe that Dinur's gap amplification step in fact already produces a CSP with this Label Cover structure, allowing us to skip the reduction step. We can thus focus on alphabet reduction when the CSP we are reducing from has the Label Cover structure, and is over a fixed, albeit large, alphabet. We then follow the influential Label Cover and Long Code framework, originally proposed in [BGS98] and strengthened in [Has01] and since then applied in numerous works on inapproximability, to reduce the CSP obtained from gap amplification, now viewed as Label Cover, to a Boolean CSP. Finally, we reduce the Boolean CSP back to a Label Cover instance (see Section 7.4) that can be plugged in as input to the gap amplification step.

Our main result is the following, which can be viewed as reproving a case of alphabet reduction from [DR06; Din07].

Theorem 107. There is a polynomial time reduction from Label Cover with soundness ϵ to a fixed template CSP with soundness δ for an absolute constant $\epsilon > 0$.

We analyze our reduction using Fourier analysis as pioneered by Has01. Usually, in this framework, we reduce from low soundness Label Cover to strong (and at times optimal) soundness of CSP. But here we start with a high soundness Label Cover, and we reduce to high soundness CSP.

We highlight a couple of differences from previous works that make our proof easier:

- We have the freedom to choose any CSP rather than trying to prove inapproximability of a CSP. We choose the following binary predicate R in our reduction: $(u; v; w; x) \in R$ if and only if $u \oplus v = w \oplus x$. This predicate appears in [Has01] in the context of proving optimal hardness for NAE-4SAT.
- In [Has01] and [BGS98], the objective is to prove optimal inapproximability, or at least to

¹While this might seem circular, as this is what the PCP reduction is trying to accomplish in the first place, the key is that this inner reduction can be highly inefficient (even triply exponential blow up is okay!), as it is applied to a constraint of constant size.

²Technically, the gap amplification step produces a version of Label Cover whose constraints are *rectangular* rather than functions, but this is a minor difference that can be easily accommodated in reductions from Label Cover.

get good soundness. However, our present goal is to prove 'just' a nontrivial soundness. (On the other hand, we also start with high soundness Label Cover.) This allows us to use a very convenient test distribution leading to a simple analysis.

- We remark that a similar statement as Theorem 107 can be also deduced using [BGS98, Section 4.1.1]. We believe that the test presented in this chapter is more direct since we benefit from ideas in [Has01].
- It is possible to perform alphabet reduction using the Hadamard code instead of the long code as described in [GO05; RS07]; the latter [RS07], similarly to our proof, avoids explicit use of Assignment Testers.
- Long code tests correspond exactly to testing whether a function is a morphism of the corresponding CSP, and as such corresponds to gadget reductions in the algebraic approach to CSP (see e.g. [BKW17]). The PCP theorem surpasses these algebraic (gadget) reductions; this is even more evident when extending the scope from CSPs to constraint satisfaction problems (PCSPs) as there are PCSPs that can be shown to be NP-hard by using PCP theorem via a natural reduction through Label Cover, but cannot be shown to be NP-hard using only algebraic reductions [AGH17; Bar+21]. In this sense, the present work shows that this strength of the PCP theorem comes from the Gap Amplification step.

Alphabet Reduction is an essential step in both the original proof of the PCP theorem as well as Dinur's proof and deserves further attention. Our proof of alphabet reduction bypasses the concept of Assignment Testers and is more intuitive in our opinion as it is just a gadget reduction. Our proof is elementary using only Parseval's identity from Fourier Analysis over the hypercube. Dinur's analysis used the Friedgut-Naor-Kalai theorem [FKN02] about Boolean functions with most of the Fourier mass at level 1. We believe that this makes our proof more accessible to readers that are new to PCPs. We also hope that this material might be useful in teaching the proof of PCP theorem as it relies only on techniques that any such basic course would cover anyway.

Outline

We start by formally defining CSP, Label Cover, and other preliminaries in Section 7.2. Then, in Section 7.3, we prove the main reduction from Label Cover to CSP. In Section 7.4, we show how the reduction can be used in the alphabet reduction step of Dinur's proof.

7.2 Preliminaries

7.2.1 Rectangular relation and the long code

In this chapter, we view the Label Cover problem as a binary CSP with the relations being projections. More generally, we consider Label Cover instances where the relations are rectangular.

Definition 108 (Rectangular relation) A relation $R \subseteq A \times B$ is said to be rectangular if there is a set C and functions $\alpha: A \rightarrow C$ and $\beta: B \rightarrow C$ such that $(a; b) \in R$ if and only if $\alpha(a) = \beta(b)$. Equivalently, R is rectangular if for all $a; a' \in A$ and $b; b' \in B$ such that $(a; b) \in R$, $(a; b') \in R$, and $(a'; b) \in R$, we have $(a'; b') \in R$.

Long code is often used in conjunction with the Label cover problem to obtain inapproximability results. Loosely speaking, the long code is the longest (error-correcting) code over the Boolean alphabet that does not repeat bits. It is constructed as follows: the long code is a Boolean code of length 2^n which encodes a value $\in \{0, 1\}^n$ into a tuple p_i whose k -th coordinate $p_{i,k} < 2^n$, is the i -th least significant digit of k in binary.

The long code can also be described in another way: we view a Boolean tuple of length 2^n as an n -ary function $f: \{0, 1\}^n \rightarrow \{0, 1\}$ (each coordinate of the tuple encodes one value). In this perspective, the code words of the long code are functions defined as $p_i(x_1; \dots; x_n) = x_i$. These functions are often called dictators.

We also remark that in the conjunction with the long-code, a rectangular constraint can be expressed as an identity. More precisely, given a rectangular relation $R \subseteq [n] \times [m]$, say $R = \{(i; j) : \alpha(i) = \beta(j)\}$ for some $\alpha: [n] \rightarrow \{0, 1\}$ and $\beta: [m] \rightarrow \{0, 1\}$, then the long codes p_i and p_j of values i, j satisfy

$$p_i(x_{(1)}; \dots; x_{(n)}) = p_j(x_{(1)}; \dots; x_{(m)})$$

for all $x_1; \dots; x_k \in \{0, 1\}$ if and only if $(i; j) \in R$. This is a key property of rectangular relations that is used implicitly in our proof.

7.2.2 Boolean Fourier analysis

As usual in Boolean Fourier analysis, we treat TRUE as -1 and FALSE as $+1$. In particular, in this notation, 'negation' is expressed as $\bar{x} = -x$, 'xor' $x \oplus y$ is expressed as $x \bar{y} + \bar{x} y$, and 'or' is expressed by the following function:

$$x \vee y = \begin{cases} 1 & \text{if } x = 1 \text{ or } y = 1, \text{ and} \\ 1 & \text{otherwise.} \end{cases}$$

We will use all the same symbols to denote the coordinatewise (or bitwise) application of these functions to tuples, e.g. $(x_1; x_2) \oplus (y_1; y_2) = (x_1 \oplus y_1; x_2 \oplus y_2)$.

We define an inner product space on functions from $\{0, 1\}^n$ to \mathbb{R} as $\langle f; g \rangle = E_x[f(x)g(x)]$.

For a set $S \subseteq [n]$, let

$$\chi_S(x_1; \dots; x_n) = \prod_{i \in S} x_i$$

The set of such functions form an orthonormal basis for all functions from $\{0, 1\}^n$ to \mathbb{R} in the above defined inner product space. Moreover, if $S \neq \emptyset$, then $E_x[\chi_S(x)] = 0$.

Definition 109 (Fourier expansion) Given a function $f: \{0, 1\}^n \rightarrow \mathbb{R}$, we can thus write it uniquely as a linear combination of this basis—

$$f = \sum_{S \subseteq [n]} \hat{f}(S) \chi_S$$

The real quantities $\hat{f}(i)$ are called the Fourier coefficients of f . We abuse the notation $\hat{f}(i)$ to denote $\hat{f}(i; g)$.

The following simple but crucial identity follows from the definitions and is all that we will need in our analysis.

Theorem 110 (Parseval's Identity) For each Boolean valued function f , i.e., $f : \{-1, +1\}^n \rightarrow \{-1, +1\}$,

$$\sum_{i \in [n]} \hat{f}(i)^2 = 1.$$

Connection to the long code

We remark, that the function $f_{i; g}$ corresponds to a valid long code: the function encoding the value i . Also observe that there is a connection between the natural distance defined by the inner product $\langle f; g \rangle$ on Boolean functions and relative Hamming distance of f and g : This is thanks to the fact that $\langle f; y \rangle = \sum_{i \in [n]} f(i) y(i)$ then $x = y$ if and only if $xy = 1$, and consequently, the relative Hamming distance of f to the long code word $\phi_i = f_{i; g}$ can be expressed as $\frac{1}{2} \sum_{i \in [n]} \hat{f}(i) = \frac{1}{2} \hat{f}(i)$. This means that the closest valid long code to a function is the ϕ_i for which $\hat{f}(i)$ is maximal.

These ideas are manifested in the common strategy in rounding a Boolean function to a long code: First make sure that coefficients $\hat{f}(i)$ for large sets are small enough, then decode to a value i that belongs to a small-enough (ideally 1-element) set with a large-enough $\hat{f}(i)$.

7.3 Label Cover to CSP

This section describes our gadget reduction from Label Cover to CSP(R) where R is the 4-ary relation over Boolean domain defined as

$$R = \{(x; x^0; z; z^0) \mid x \in \{-1, +1\}, z \in \{-1, +1\}\}.$$

Theorem 111. There exists absolute constants ϵ, δ such that given a Label Cover instance (not necessarily bipartite) $G = (V; E; \{C_e\}_{e \in E})$ with rectangular constraints, there is a reduction from G that outputs an instance of CSP(R) such that $\text{size}(I) = O(\text{size}(G))$ and

- If G is satisfiable, then I is satisfiable as well.
- If no labeling can satisfy ϵ fraction of constraints of G , then no assignment can satisfy δ fraction of constraints in I for all ϵ .

Since the domain of CSP(R) is Boolean, the above reduces from an alphabet of arbitrary size to the alphabet of size 2. We note that the constant $\delta(\epsilon)$ above hides exponential dependency on ϵ .

We describe the reduction as a probabilistic checker of a solution ϕ to G encoded using a long code, i.e., the proof contains for each $u \in V$ a word $f_u : \{-1, +1\}^j \rightarrow \{-1, +1\}$. In other words, we design the test in such a way that if ϕ is a solution to G , then the assignment $f_u = \phi_{S(u)}$ passes the test. This will then immediately give the completeness of the reduction.

The test is as follows: Sample an edge $e = (u; v)$ from E uniformly at random, and then with equal probability do one of the following

1. run a long code test inside e ;
2. run a long code test inside e ;
3. run a constraint test between u and v .

We describe the long code test and the constraint test below. Both query the respective tables of f_u and f_v at some 4 bits that are generated by a certain randomized algorithm, and then check whether these 4 Boolean values satisfy the predicate defined above.

This checker can be viewed as a gadget reduction in the following sense: We replace each vertex $u \in V$ with 2^{j-1} Boolean variables labeled $f_u(x)$ for $x \in \{0, 1\}^{j-1}$ (we see an assignment to such variables as a function $f : \{0, 1\}^{j-1} \rightarrow \{0, 1\}$), and each edge $e = (u; v)$ with a set of weighted 4-ary constraints Γ_e and f_v , each involving the relation R and some values of f_u and f_v (the result is therefore an instance of CSP(R)). These constraints depend only on the relation R .

To simplify some notation, we assume n is even. We also assume that the tables f_u are folded so $f_u(x) = f_u(\bar{x})$. This is a standard technique. Such a folding is ensured by including only one variable of each pair x , and if the test queries f_u at the bit corresponding to some x that is not included, the bit $f_u(\bar{x})$ is queried instead, and the value is negated. As a consequence of this folding, we have to allow for negation of variables. An important and useful consequence of this is that all 'even' Fourier coefficients vanish, i.e., $\hat{f}(S) = 0$ for all S such that $|S|$ is even. We remark that folding can be avoided in the construction of the gadget. Nevertheless, it considerably simplifies the calculations below. Further, for calculations, it is useful to view R as a predicate: $f : \{0, 1\}^4 \rightarrow \{0, 1\}$ defined as

$$f(x; x^0; z; z^0) = 1 - \frac{(xx^0 + 1)(zz^0 + 1)}{4}.$$

It is easy to check that $f(x; x^0; z; z^0) = 1$ if and only if $(x; x^0; z; z^0) \in R$.

Let us now describe the two probabilistic checkers.

7.3.1 Long code test

The long code test has on input a table of a function f (either f_u or f_v), and it is supposed to check whether this function is a code word of the long code, i.e., there is $p \in \{0, 1\}^n$ such that $f = p$. We design the test so that all these words pass with probability 1. Since we are only using the predicate f , this further limits possible checks. In fact, we include all checks of the form $(x_1; x_2; x_3; x_4) \in R$ that are passed by all dictators.

Long code test. Given $f : \{0, 1\}^n \rightarrow \{0, 1\}$ to test against being a long code word. Choose $x; y; z; \bar{x}; \bar{y}; \bar{z} \in \{0, 1\}^n$ uniformly at random. Test whether

$$(f(x); f(\bar{x} \oplus y); f(z); f(\bar{z} \oplus y)) \in R \tag{7.1}$$

³Any function that passes any such test with probability 1 is called a *dictator* or *dictator function*, see also [BKW17].

Note that for all $x, y, z \in \mathbb{F}_2^k$, $g(x, x \oplus y); z; z \oplus y) \in \mathbb{R}$: This implies that any dictator function passes the test with probability 1, and therefore provides the completeness of the test. We also note that this test can give some false positives, e.g. the function $\neg p_i(x)$ passes the test with probability 1, but is not a long code word. It is in fact a negation of the word p_i . It can be checked that all functions that pass are either long code words, or their negations. In the decoding, we simply decode the above function.

The following lemma bounds the probability that the test accepts in the means of the Fourier coefficients. We remark, that since we want to ensure that f is as close to a valid long code as possible, the probability should decrease as the coefficients for ϵ increase. Indeed, the lemma states that this is the case.

Lemma 112. Assuming that f is folded, the probability the long code test accepts is at most

$$1 - \frac{3}{16} \sum_{|j|>1} \hat{f}(j)^2$$

Proof. Assume $f(x) = \sum_j \hat{f}(j) \chi_j(x)$. The probability that the test accepts is

$$\begin{aligned} & \mathbb{E}_{x,y,z} (f(x); f(x \oplus y); f(z); f(z \oplus y)) \quad \# \quad (7.2) \\ &= \mathbb{E}_{x,y,z} \frac{1}{4} (f(x)f(x \oplus y) + 1 - f(z)f(z \oplus y) + 1) \\ &= \frac{3}{4} \left(\frac{1}{4} \mathbb{E}_{x,y} f(x)f(x \oplus y) - \frac{1}{4} \mathbb{E}_{y,z} f(z)f(z \oplus y) \right) \\ & \quad + \frac{1}{4} \mathbb{E}_{x,y,z} f(x)f(x \oplus y)f(z)f(z \oplus y) \end{aligned}$$

We further simplify this expression one term at a time.

$$\begin{aligned} \mathbb{E}_{x,y} f(x)f(x \oplus y) &= \mathbb{E}_{x,y} \sum_j \hat{f}(j) \hat{f}(j \oplus x) \chi_j(x) \chi_{j \oplus x}(x \oplus y) \quad (7.3) \\ &= \sum_j \hat{f}(j) \hat{f}(j \oplus x) \mathbb{E}_x [\chi_j(x) \chi_{j \oplus x}(x)] \mathbb{E}_y [\chi_{j \oplus x}(x \oplus y)] \\ &= \sum_j \hat{f}(j)^2 \mathbb{E}_y [\chi_{j \oplus x}(x \oplus y)] = \sum_j \hat{f}(j)^2 (1/2)^{|j|} \end{aligned}$$

The third equality follows since χ_j and $\chi_{j \oplus x}$ are orthogonal if $j \neq 0$. The last equality follows from the fact that $\mathbb{E}_y [\chi_{j \oplus x}(x \oplus y)] = (1/2)^{|j|}$ and coordinates are chosen independently. Similarly, we get that

$$\mathbb{E}_{y,z} f(z)f(z \oplus y) = \sum_j \hat{f}(j)^2 (1/2)^{|j|} \quad (7.4)$$

Moving to the next term, we get

$$\begin{aligned} & \mathbb{E}_{x,y,z} f(x)f(x \oplus y)f(z)f(z \oplus y) \\ &= \sum_j \hat{f}(j)^2 \hat{f}(j \oplus x)^2 \mathbb{E}_y [\chi_{j \oplus x}(x \oplus y) \chi_{j \oplus x}(x \oplus y)] = \sum_j \hat{f}(j)^2 \hat{f}(j \oplus x)^2 (1/2)^{2|j|} \quad (7.5) \end{aligned}$$

The last equality follows since $E_{y_i} [(x_i - y_i) - (x_i - y_i)] = E_{y_i} [0] = 0$ and $E_{y_i} [(x_i - y_i) - (x_i - y_i)] = E_{y_i} [0] = 0$. The overall acceptance probability is then

$$\begin{aligned} & \frac{3}{4} \sum_{j=1}^X f(\omega_j)^2 (1-2)^j + \frac{1}{4} \sum_{j=1}^X f(\omega_j)^2 f(\omega_{j+1})^2 (1-2)^{j+1} \\ &= 1 - \frac{1}{2} \sum_{j=1}^X f(\omega_j)^2 (1-2)^j + 1 - \frac{1}{4} \sum_{j=1}^X f(\omega_j)^2 f(\omega_{j+1})^2 (1-2)^{j+1} \end{aligned} \quad (7.6)$$

where for the last equality, we used Parseval's identity. Further, we assumed $f(\omega_j) = 0$ for all j such that j is even. Restricting the sums to j with odd cardinality, and using that for such disjoint $j, j+1$ is even, the last expression of (7.6) can be bounded from above by

$$1 - \frac{1}{2} \sum_{j>1} f(\omega_j)^2 (1-2)^j \leq 1 - \frac{1}{4} \sum_{j>1} f(\omega_j)^2 f(\omega_{j+1})^2 (1-2)^{j+1} \leq 1 - \frac{3}{16} \sum_{j>1} f(\omega_j)^2 \quad (7.7)$$

which concludes the proof. □

7.3.2 Constraint test

The constraint test has on input tables for functions f and g corresponding to some u and v such that $(u; v) \in E$, and it is supposed to test (assuming f and g are correct long code words) whether the values these functions encode satisfy the constraint given by a rectangular relation. We construct the test in a similar way as the long code test: We test functions f and g in 4 bits in such a way that long code words encoding satisfying values pass. In contrast with the long code test, we do not include all such tests, but only a selection; in particular, we include only tests that query two values from each function.

We assume that the constraint relation is given by $R: [n] \times [m] \rightarrow \{0, 1\}$ such that $(i; j) \in R$ if and only if $(i) = (j)$, and we denote by \mathbf{y} the vector in $\{0, 1\}^n$ such that $y(i) = (i)$.

Constraint test. Given $f; g: \{0, 1\}^n \rightarrow \{0, 1\}$ to test against satisfying a constraint given by $(i; j) \in R$ if and only if $(i) = (j)$ for $x \in \{0, 1\}^n; z \in \{0, 1\}^m$. Choose $x; z \in \{0, 1\}^n$ and $y \in \{0, 1\}^m$ uniformly at random, and test whether

$$(f(x); f(x \oplus y); g(z); g(z \oplus (y))) \in R \quad (7.8)$$

Note that if both f and g are dictators, say $f = p_i$ and $g = p_j$, such that $(i) = (j) = k$ then the above test accepts with probability 1. Indeed, the tuple gets evaluated to

$$(x(i); (x \oplus y)(i); z(j); (z \oplus (y))(j)) = (x(i); x(i) \oplus y(k); z(j); z(j) \oplus y(k))$$

which is in R for all $x; y$ and z . This provides the completeness of the test.

In the analysis below, we will use the following notation.

Definition 113. Let $\mathcal{C} \subseteq \{0,1\}^n$ and $\mathcal{D} \subseteq \{0,1\}^m$, we denote by $\mathcal{C} \times \mathcal{D}$ the subset of $\{0,1\}^{n+m}$ defined by $\mathcal{C} \times \mathcal{D} = \{x : x_{(1)} \in \mathcal{C}, x_{(2)} \in \mathcal{D}\}$.

The goal of the constraint check is to ensure that functions which are far from valid long codes that encode values satisfying the constraint pass with low probability. Unfortunately, the test gives a lot of false positives: it accepts any pair of functions f and g such that $\mathcal{C} \times \mathcal{D} = \mathcal{C} \times \mathcal{D}$ with probability $1/4$. This is nevertheless good-enough since the long code test provides that relevant and \mathcal{C} contain only one element, and $\mathcal{C} \times \mathcal{D} = \mathcal{C} \times \mathcal{D}$ if and only if $\mathcal{C} = \emptyset$.

Naturally, the pairs of f and g with $\mathcal{C} \times \mathcal{D} = \mathcal{C} \times \mathcal{D}$ will appear in the analysis below. A useful fact that will simplify the computation below is that $\sum_{i \in \mathcal{C}} x_{(i)} = \sum_{i \in \mathcal{C}} x_i$, for all $x \in \{0,1\}^{n+m}$, which implies that

$$\sum_{i \in \mathcal{C}} x_{(i)} = \sum_{i \in \mathcal{C}} x_i$$

Lemma 114. Given that both f and g are folded, the probability that the consistency test accepts is at most

$$1 - \frac{1}{4} \sum_{i,j \in \mathcal{C}} \hat{f}(i)^2 \hat{g}(j)^2$$

Proof. We can compute the acceptance probability in the same way as before, i.e., as

$$\begin{aligned} \frac{3}{4} &= \frac{1}{4} \mathbb{E}_{x,y} [f(x)f(x \oplus y)] + \frac{1}{4} \mathbb{E}_{z,y} [g(z)g(z \oplus y)] \\ &= \frac{1}{4} \mathbb{E}_{x,y,z} [f(x)f(x \oplus y)g(z)g(z \oplus y)] \quad (7.9) \end{aligned}$$

We have

$$\mathbb{E}_{x,y} [f(x)f(x \oplus y)] = \sum_{i \in \mathcal{C}} \hat{f}(i)^2 \mathbb{E}_y [f(i \oplus y)] = \sum_{i \in \mathcal{C}} \hat{f}(i)^2 \mathbb{E}_y [f(i)] = 0 \quad (7.10)$$

where the last equality holds since \mathcal{C} is odd, and consequently $\mathbb{E}_y [f(i \oplus y)] = 0$. Similarly, $\mathbb{E}_{x,z} [g(z)g(z \oplus y)] = 0$.

⁴We note that $\mathcal{C} \times \mathcal{D} = \mathcal{C} \times \mathcal{D}$ is equivalent to $\sum_{i \in \mathcal{C}} x_{(i)} = \sum_{i \in \mathcal{C}} x_i$ for all $x \in \{0,1\}^{n+m}$.

$(: y)]$ vanishes. Thus the probability that the test accepts is

$$\begin{aligned}
 & \frac{3}{4} - \frac{1}{4} E_{x;y;z} f(x)f(x-y)g(z)g(z-(y)) \tag{7.11} \\
 &= \frac{3}{4} - \frac{1}{4} \sum_{j>1} \hat{f}(j)^2 \hat{g}(j)^2 E_y [(y) - (y)] \\
 &= \frac{3}{4} + \frac{1}{4} \sum_{j>1} \hat{f}(j)^2 \hat{g}(j)^2 E_y [(y) - (y)] \\
 &= \frac{3}{4} + \frac{1}{4} \sum_{j>1} \hat{f}(j)^2 \hat{g}(j)^2 E_y [\chi_j(y) - \chi_j(y)] \\
 &= \frac{3}{4} + \frac{1}{4} \sum_{j>1} \hat{f}(j)^2 \hat{g}(j)^2 \\
 &= 1 - \frac{1}{4} \sum_{j>1} \hat{f}(j)^2 \hat{g}(j)^2
 \end{aligned}$$

where the second equality follows from j being odd, and the last equality follows from the Parseval's identity. Since $(i) = (j)$ implies that $[f \hat{g}] = [f \hat{g}]$, the claim follows. \square

7.3.3 The full test

Putting the analysis of the two tests together we get the following.

Lemma 115. Given that both f and g are folded, the probability that the joint test accepts is at most

$$1 - \frac{1}{16} \sum_{j>1} \hat{f}(j)^2 + \sum_{j>1} \hat{g}(j)^2 + \sum_{i,j : (i) \neq (j)} \hat{f}(i)^2 \hat{g}(j)^2$$

Proof. Follows directly from Lemmas 112 and 114. \square

Finally, we are ready to prove the main theorem of this section.

Proof of Theorem 111 The completeness follows in a straightforward way from the two comments after the description of the tests. We prove the soundness. Suppose that the test passes with probability $1 - \epsilon$. We will show that this implies that there is an assignment to the Label Cover instance that satisfies $(1 - 16\epsilon)$ -fraction of constraints.

Our decoding is as follows: for a node $v \in V$, decode v to $i \in \mathbb{Z}_2$ with probability proportional to $\hat{f}_v(i)^2$. Intuitively, we decode to the value with higher probability iff v is closer to the code word $p_i = f_i \hat{g}$ or its negative $-p_i$ (see also Section 7.2.2). We will show that in expectation, this decoding satisfies at least $1 - 16\epsilon$ fraction of constraints, which proves that there exists a labeling that satisfies at least $1 - 16\epsilon$ fraction of constraints.

Let $1 - \epsilon_e$ denote the probability that the test passes when we pick edges. If the test passes with probability $1 - \epsilon$, we know that $E_e[\epsilon_e] = \epsilon$. Suppose that we pick $e = (u; v)$ with f_u and g_v being

the functions corresponding to ϕ and ψ respectively. From the above lemma, we have that

$$1 - \epsilon \leq 1 - \frac{1}{16} \sum_{j>1} \phi(j)^2 + \sum_{j>1} \psi(j)^2 + \sum_{ij: (i)\in(j)} \phi(i)^2 \psi(j)^2; \quad (7.12)$$

and therefore,

$$16\epsilon \leq \sum_{j>1} \phi(j)^2 + \sum_{j>1} \psi(j)^2 + \sum_{ij: (i)\in(j)} \phi(i)^2 \psi(j)^2; \quad (7.13)$$

The probability that our decoding satisfies edge of Label Cover is at least

$$\begin{aligned} \sum_{ij: (i)\in(j)} \phi(i)^2 \psi(j)^2 &= 1 - \sum_{ij: (i)\in(j)} \phi(j)^2 \psi(i)^2 \\ &= 1 - \sum_{\substack{j>1, \text{ or } j>1, \text{ or } =fj \text{ and } =fjg \text{ and } (i)\in(j)}} \phi(j)^2 \psi(i)^2 - \sum_{\substack{j>1 \\ (i)\in(j)}} \phi(i)^2 \psi(j)^2 \\ &= 1 - \sum_{j>1} \phi(j)^2 - \sum_{j>1} \psi(j)^2 - \sum_{ij: (i)\in(j)} \phi(i)^2 \psi(j)^2 \\ &= 1 - 16\epsilon \end{aligned} \quad (7.14)$$

where the first equality follows from Parseval's identity. Thus, the expected number of constraints satisfied by the labeling is at least $\mathbb{E}[\sum_{ij: (i)\in(j)} \phi(i)^2 \psi(j)^2] = 1 - 16\epsilon$ which proves the required claim with $C = 1/16$. \square

Theorem 107 is stated without the assumption that the constraints are rectangular. This slightly more general version follows from Theorem 111 by a standard reduction which we describe below, in the proof of Theorem 116.

7.4 CSP to Label Cover

In this section, we recall the basic structure of Dinur's proof of PCP Theorem, and show how the previous reduction can be used in the alphabet reduction step of Dinur's proof. The resulting proof requires a gap amplification step for which we refer to Dinur's paper [Din07].

We first prove that the previous reduction can be combined with standard reductions to get back Label Cover from the CSP.

Theorem 116 (Alphabet reduction) Given a Label Cover instance $\mathcal{G} = (V; E; \{ \sigma_{ij} \}_{(i,j) \in E})$ with rectangular constraints, there is a polynomial time reduction that outputs another Label Cover instance with rectangular constraints $\mathcal{G}^0 = (V^0; E^0; \{ \sigma_{ij}^0 \}_{(i,j) \in E^0})$ with alphabet size such that $|V^0| \leq |V|$ is an absolute constant, $\text{size}(\mathcal{G}^0) = O(\text{size}(\mathcal{G}))$ and

- If \mathcal{G} is satisfiable, then \mathcal{G}^0 is satisfiable as well.
- If every labeling violates fraction of constraints ϵ of \mathcal{G} , then every labeling violates fraction of constraints ϵ^0 for an absolute constant C .

Proof. We first convert the Label Cover instance G to a CSP instance as in Theorem 111. The CSP instance can be converted to bipartite Label Cover using standard clause-variable Label-coverization technique. We include the proof here for the sake of completeness. We have n vertices x_1, x_2, \dots, x_n corresponding to the variables on the left L , and m vertices corresponding to constraints C_1, C_2, \dots, C_m of I on the right R . The label set is binary on the left, and satisfying assignments (at most 16) on the right corresponding to the possible assignments to four variables in the constraint. We add an edge between $u \in L$ and $v \in R$ if $x_u \in C_v$. The constraint on this edge enforces that the assignment α is consistent with the assignment α_v that assigns to x_u .

If there is a satisfying labeling α of G , there is a satisfying assignment α . Using this, we can assign the variables on the left the satisfying assignment, and the corresponding assignment to tuples for the vertices of constraints on the right, and thus get a satisfying assignment α to G . Suppose that every labeling violates at least a fraction of constraints ϵ . From Theorem 111, every assignment violates at least ϵ fraction of constraints in I . Suppose there is a labeling to G^0 that satisfies ϵ^0 fraction of constraints. Consider the assignment obtained by this labeling on the left. This assignment violates at least ϵ fraction of constraints in I . Note that this should violate at least $\epsilon^0 m$ constraints in G^0 and thus $\epsilon^0 \leq \epsilon m$ for an absolute constant ϵ^0 . The constraints are in fact projections, and thus are rectangular too. \square

In order for us to use this as Composition step in the PCP of Dinur, we need the final observation that the output of Gap Amplification applied to a CSP with rectangular constraints results in a Label Cover with rectangular constraints. Dinur achieves gap amplification by 'graph powering' which is described more formally below.

Definition 117 (Constraint Graph Powering) Given a d -regular Label Cover (a.k.a. Constraint graph) $G = (V; E; \dots)$, we obtain t -th power of it $G^t = (V; E^0, \dots)$ as follows:

- Vertices. The vertices in G^t are the same as vertices in G .
- Edges. u and v are connected by t parallel edges in E^0 if there are exactly t paths of length t between u and v in G .
- Alphabet. The alphabet of G^t is $\Sigma^{d^{dt=2e}}$. A value $a \in \Sigma^{d^{dt=2e}}$ is interpreted as assigning values a_u to $d^{dt=2e}$ elements from Σ . This value is treated as a 's opinion on (u) , the set of all vertices with $dt=2e$ distance from u .
- Constraints. An edge $(u; v) \in E^0$ is satisfied by $(a; b) \in \Sigma^{d^{dt=2e}}$ if and only if the following holds: there is an assignment $(u) [(v)]$ that satisfies every constraint (e) where $e \in E \setminus ((u) \cup (v))$, and such that

$$8u^0 \in (u); (u) = a_{u^0}; 8v^0 \in (v); (v) = b_{v^0}$$

where a_{u^0} (and respectively b_{v^0}) is the value a (and resp b) assigned to u^0 (and resp v^0).

The output G^t is also a binary CSP, and hence can be viewed as a Label Cover. We claim that if every constraint of G is rectangular, then every constraint of G^t is rectangular as well. Let $e = (u; v)$ be an edge in E^0 with constraint relation R_e . Suppose $(a; b); (a^0; b^0); (a; b^0) \in R_e$. This implies that for all $(u^0; v^0) \in E \setminus ((u) \cup (v))$ with constraint relation R_e ,

$$(a_{u^0}; b_{v^0}); (a_{u^0}^0; b_{v^0}^0); (a_{u^0}; b_{v^0}^0) \in R_{e^0}$$

Since R_{ce} is rectangular, $(a_u^0, b_v^0) \in R_{ce}$ as well. As this holds for all such u and v , $(a^0, b^0) \in R_e$, thus proving that R_e is a rectangular relation.

Combined with the preprocessing step, the gap amplification theorem of Dinur can be rewritten as follows.

Theorem 118 (Gap amplification) Fix a parameter t . Given a Label Cover $G = (V; E; \sigma; \tau)$ where σ is an absolute constant, there is a polynomial time reduction to output a rectangular Label Cover instance $G^0 = (V^0; E^0; \sigma^0; \tau^0)$ with the alphabet size $|V^0| = \alpha(j; t)$ such that

- If G is satisfiable, G^0 is satisfiable as well.
- If every labeling violates at least fraction ϵ of the constraints of G , then every labeling violates at least (ϵ^t) fraction of the constraints of G^0 .

Choosing t large enough constant and iterating Theorem 116 and Theorem 118 $\log(1/\epsilon)$ times proves the PCP theorem.

7.5 Derandomization of the gadget decoding

In this appendix, we provide a derandomization of the decoding used in Theorem 111. This requires only a little additional argument. The idea is, instead of decoding with probability $\hat{f}(i)^2$, to decode to the i with the largest $\hat{f}(i)^2$. We set ϵ_f to be such. In this light, the reduction can be analyzed by analyzing completeness and soundness of the gadget separately without considering the rest of the instance. The following lemma then shows that the gadget has perfect completeness and soundness ϵ_{99} not depending on the parameters ϵ and δ (the alphabet sizes), and ϵ_f .

Lemma 119. There is a gadget with inputs $s, m, k, \sigma : [n] \rightarrow [k]$, and $\tau : [m] \rightarrow [k]$ that produces an instance of CSP(R) with variables $f = (a_1; \dots; a_n)$ and $g = (a_1; \dots; a_m)$ such that

1. if $(i) = (j)$ then p_i and p_j satisfy all the constraints, and
2. if at least 99% of the constraints are satisfied, then $(i_f) = (i_g)$.

Proof. First, we bound the probability that the test accepts. For the long code test, starting with the first expression in (7.7), we obtain the following bound.

$$1 - \frac{1}{2} \sum_{j > 1} \hat{f}(j)^2 \leq \frac{1}{4} \sum_{i=j} \hat{f}(i)^2 \hat{f}(j)^2 (1 - 2^{-j})$$

$$1 - \frac{3}{16} \sum_{j > 1} \hat{f}(j)^2 \leq \frac{1}{16} \sum_{i \in j} \hat{f}(i)^2 \hat{f}(j)^2$$

For the consistency test, we use the bound from Lemma 114. Thus the overall probability that the

whole test accepts is at most

$$1 - \frac{1}{16} \sum_{j>1} \hat{f}(j)^2 - \frac{1}{48} \sum_{i \in j} \hat{f}(i)^2 \hat{f}(j)^2 - \frac{1}{16} \sum_{j>1} \hat{g}(j)^2 - \frac{1}{48} \sum_{i \in j} \hat{g}(i)^2 \hat{g}(j)^2 - \frac{1}{12} \sum_{ij: (i) \in (j)} \hat{f}(i)^2 \hat{g}(j)^2$$

Given that the acceptance probability is at least $\frac{2}{3} > 1 - \frac{1}{96}$, we get that

$$\sum_{j>1} \hat{f}(j)^2 \leq \frac{1}{6} \tag{7.15}$$

$$\sum_{i \in j} \hat{f}(i)^2 \hat{f}(j)^2 \leq \frac{1}{2} \tag{7.16}$$

$$\sum_{i \in j} \hat{g}(j)^2 \leq \frac{1}{6} \tag{7.17}$$

$$\sum_{j>1} \hat{g}(i)^2 \hat{g}(j)^2 \leq \frac{1}{2} \tag{7.18}$$

$$\sum_{ij: (i) \in (j)} \hat{f}(i)^2 \hat{g}(j)^2 \leq \frac{1}{8} \tag{7.19}$$

From Parseval's identity and (7.15), we get that $\sum_i \hat{f}(i)^2 = \frac{5}{6}$. Recall that f is such that $\hat{f}(i)^2$ is maximal. Then using the above and (7.16), we obtain that

$$\sum_i \hat{f}(i_f)^2 = \sum_i \hat{f}(i)^2 \sum_i \hat{f}(i)^4 = \sum_{ij} \hat{f}(i)^2 \hat{f}(j)^2 \sum_{i \in j} \hat{f}(i)^2 \hat{f}(j)^2 = \left(\frac{5}{6}\right)^2 \cdot \frac{1}{2} = \frac{4}{9} \tag{7.20}$$

Similarly, from (7.17) and (7.18), we get $\sum_i \hat{g}(i_g)^2 = \frac{4}{9}$. Finally, since $\sum_i \hat{f}(i_f)^2 \hat{g}(i_g)^2 = \left(\frac{4}{9}\right)^2 > \frac{1}{8}$, we have that $(i_f) = (i_g)$ otherwise (7.19) cannot be true. \square

Theorem 111 can be also directly obtained from this lemma albeit with a worse constant than in the above proof: Let $\epsilon = \frac{1}{96}$ and assume that $\epsilon < 1$. Given that the resulting CSP instance has an assignment fails no more than ϵ -fraction of the constraints, we derive that in at least $(1 - \epsilon)$ -fraction of the gadgets, no more than ϵ -fraction of constraints are unsatisfied. Lemma 119 then shows that the assignment $\sigma|_{i_{f_u}}$ is an assignment of the Label Cover instance that satisfies all the constraints corresponding to these gadgets. This completes the proof.

Part II

Structured instances

Chapter 8

Multidimensional Packing and Scheduling Problems

8.1 Introduction

Bin Packing and Multiprocessor Scheduling (also known as Makespan Minimization) are some 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 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. These problems are well understood in terms of approximation algorithms: both the problems are NP-hard, and have a Polynomial Time Approximation Scheme (PTAS) [VL81; HS87].

In this chapter, we study the approximability of the multidimensional generalizations of these problems. The corresponding problems are Vector Bin Packing and Vector Scheduling. Apart from their theoretical importance, these problems are widely applicable in practice [Spi94; ST12; Pan+11] 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 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 asymptotic PTAS for 2-dimensional Vector Bin Packing, assuming $P \neq NP$. On the algorithmic front, the PTAS for Bin Packing [VL81] easily implies a $(1 + \epsilon)$ approximation for Vector Bin Packing. When a part of the input, this

¹Very recently, Ray [Ray21] found an oversight in Woeginger's original proof and gave a revised APX hardness proof for the problem.

²The asymptotic approximation ratio (formally defined in Section 8.2) 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 chapter for Vector Bin Packing are asymptotic.

is almost tight: there is a lower bound³ shown by [CK04]. When d is a fixed constant, much better algorithms are known [CK04; BCS09; BEK16] that get a $(1 + O(1))$ 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 $(1 + o(1))$ 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 [Chr+17]. It also appeared in a recent report by Bansal [Ban17] on open problems in scheduling. In fact, to the best of our knowledge, super constant integrality gap instances for the configuration LP relaxation of the problem were also not known. For the integer instances i.e. when the vectors are $\{0, 1\}^d$ (which are the hard instances for Vector Scheduling and Vector Bin Covering), there is an asymptotic PTAS since there are $O_d(1)$ item types.

In the Vector Scheduling problem, given a set of vector jobs in $[0, 1]^d$, and m identical machines, the objective is to assign the jobs to machines to minimize the maximum 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 with a fixed constant. When d is part of the input, they obtained a $(\log^2 d)$ factor approximation algorithm. They also showed that it is NP-hard to obtain a $(1 + \epsilon)$ factor approximation algorithm for the problem, for any constant ϵ . Meyerson, Roytman, and Tagiku [MRT13] gave an improved $(\log d)$ factor algorithm while the current best factor is $\frac{\log d}{\log \log d}$ due to Harris and Srinivasan [HS19] and Im, Kell, Kulkarni, and Panigrahi [Im+19]. 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 machines setting.

8.1.1 Our Results

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

Vector Bin Packing

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

Theorem 120. There exists an integer d_0 and a constant $\epsilon > 0$ such that for all constants $d \geq d_0$, d -dimensional Vector Bin Packing has no asymptotic $(1 + \epsilon)$ factor polynomial time approximation algorithm unless $\mathcal{P} = \text{NP}$.

³[CK04] actually gave a $\frac{1}{2}$ hardness, but it has been shown later (see e.g., [Chr+17]) that a slight modification of their reduction gives a $\frac{1}{2}$ hardness.

⁴The algorithms are now allowed to run in time $\epsilon^{f(d)}$, for some function f .

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} 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 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

$$\sum_{v \in S} f(v) \leq \mathbf{1}$$

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 + ϵ) n set cover hardness [Fei98] have a large packing dimension that grows with n , which we cannot afford as we are operating in the constant 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 simple⁵ i.e., every pair of sets intersect in at most one element.
2. The cardinality of each set is at most a fixed constant.
3. Each element in the family is present in at most $k^{O(1)}$ 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 (log k). We observe that by modifying the parameters slightly in their proof, we can obtain hardness of simple bounded set cover.

We prove that simple bounded set systems have packing dimension at most $(\log k)$. Thus, the $(\log k)$ simple bounded set cover hardness translates to 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 $k^{O(1)}$ structured simple bounded set systems on the same universe, and then we give an embedding (using $k^{O(1)}$ 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

⁵Simple set families are also known as linear set families.

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.

Vector Scheduling

For the Vector Scheduling problem, we obtain $(\log d)^1$ hardness under $\text{NP}^* \text{ZPTIME } n^{(\log n)^{O(1)}}$, almost matching the $\Theta\left(\frac{\log d}{\log \log d}\right)$ algorithms [HS19; Im+19].

Theorem 121. For every constant $\epsilon > 0$, assuming $\text{NP}^* \text{ZPTIME } n^{(\log n)^{O(1)}}$, d -dimensional Vector Scheduling has no polynomial time $(\log d)^{1-\epsilon}$ -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 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 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 [Im+19] in the context of proving lower bounds for online Vector Scheduling. It was also used implicitly in the $(1-\epsilon)$ 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 m^1 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 $B = (\log n)^\epsilon$, for some constant $\epsilon > 0$ under the assumption that $\text{NP}^* \text{ZPTIME } n^{(\log n)^{O(1)}}$.

We then amplify this hardness $B = (\log n)^\epsilon$, for every constant $\epsilon > 0$. 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 when 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 $B = (\log n)^\epsilon$ for any constant ϵ . We then use lexicographic product of graphs to amplify this result into the hardness of original Monochromatic Clique problem with $B = (\log n)^\epsilon$ for any constant ϵ under the same assumption $\text{NP}^* \text{ZPTIME } n^{(\log n)^{O(1)}}$. 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 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 122. For every constant $\epsilon > 0$, d -dimensional Vector Scheduling is NP-hard to approximate within $(\log \log d)^\epsilon$ when d is part of the input.

Consider the case when each vector job is from \mathbb{R}^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 an 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 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. 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 123. There exists a constant $\epsilon > 0$ such that assuming $\text{NP} \subseteq \text{DTIME}(n^{\epsilon \log \log n})$, d -dimensional Vector Scheduling is hard to approximate within $(\log d)^\epsilon$ when d is part of the input.

8.1.2 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 [Gar+76] gives $O(d)$ competitive ratio, and Azar, Cohen, Kamara, and Shepherd [Aza+13] recently gave an almost matching $(d^{1-\epsilon})$ lower bound. For the d -dimensional Vector Scheduling, Im, Kell, Kulkarni, and Panigrahi [Im+19] gave a $\frac{\log d}{\log \log d}$ competitive online algorithm and proved a matching lower bound.

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 d -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 [BK14a] 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 Chleb and Chlebová [CC06]. We refer the reader to the excellent survey [Chr+17] regarding the geometric problems.

8.1.3 Organization

We first define the multidimensional problems and the Label Cover problem formally in Section 8.2. Next, we prove the hardness results for Vector Bin Packing and Vector Scheduling in Section 8.3, Section 8.4 respectively.

8.2 Preliminaries

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

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

Definition 124. (Vector Bin Packing) In the Vector Bin Packing problem, the input is a set of 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

$$\sum_{j \in A_i} v_j \leq \mathbb{1} \quad \forall i \in [m]$$

Definition 125. (Vector Scheduling) In the Vector Scheduling problem, the input is a set of 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 load on a machine.

$$\max_{i \in [m]} \sum_{j \in A_i} v_j$$

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 [Chr+17] of asymptotic approximation ratio of an algorithm A for the Vector Bin Packing problem.

Definition 126. (Asymptotic Approximation Ratio) The asymptotic approximation ratio of an algorithm A for the Vector Bin Packing problem is

$$\rho_A = \limsup_{n \rightarrow \infty} \frac{A(I)}{\text{OPT}(I)} : \text{OPT}(I) = n$$

where I denotes the set of all possible Vector Bin Packing instances.

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

Label Cover. We define the Label Cover problem:

Definition 127. (Label Cover) In an instance of the Label Cover problem $G = (V = L \cup R; E; \ell_L; \ell_R)$ with $|L| = |R|$, the input is a bipartite graph $G = (L \cup R; E)$ 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 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 128. 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 [Aro+98] combined with Raz's parallel repetition [Raz98] yields the following strong inapproximability of Label Cover problem.

Theorem 129. There exists an absolute constant $\epsilon > 0$ such that for every integer n and $\delta > 0$, there is a reduction from 3-SAT instance ϕ over n variables to Label Cover instance $G = (V = L \cup R; E; \ell_L; \ell_R)$ with $|V| \leq n^{O(\log(1/\delta))}$, $|L| = |R| \leq \frac{1}{\delta}^c$ satisfying the following:

1. (Completeness.) If ϕ is satisfiable, there exists a labeling σ that satisfies all the constraints.
2. (Soundness.) If ϕ is not satisfiable, no labeling can satisfy a fraction of the constraints of G .
3. (Biregularity.) The graph $G = (L \cup R; E)$ is biregular with degrees on either side bounded by $\text{poly}(1/\delta)$.

Furthermore, the running time of the reduction is $\text{poly}(n; 1/\delta)$.

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

Theorem 130. There exist absolute constants $\epsilon > 0$ such that for every n and $\delta > 0$, there is a reduction from 3-SAT instance ϕ over n variables to Label Cover instance $G = (V = L \cup R; E; \ell_L; \ell_R)$ with $|V| \leq n^{1+o(1) \cdot \frac{1}{\delta}^c}$, $|L| = |R| \leq 2^{(\frac{1}{\delta})^{c_0}}$ satisfying the following:

1. (Completeness.) If ϕ is satisfiable, there exists a labeling σ that satisfies all the constraints.

2. (Soundness.) If is not satisfiable, no labeling can satisfy a fraction of the constraints of G .

3. (Biregularity.) The graph $H = (R; E)$ is biregular with degrees on either side $\leq \frac{1}{\epsilon}$.

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

8.3 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 120.

8.3.1 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 131. 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

$$\|f(S)\|_1 \leq k.$$

If no such embedding exists, we say $\text{pdim}(\mathcal{S})$ 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 132. 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,

$$\text{pdim}(\mathcal{S}_1 \cup \mathcal{S}_2) = \text{pdim}(\mathcal{S}_1) + \text{pdim}(\mathcal{S}_2)$$

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 $S \in \mathcal{S}_1$,

$$\|f_1(S)\|_1 \leq k_1$$

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 $S \in \mathcal{S}_2$,

$$\|f_2(S)\|_1 \leq k_2$$

if and only if $S \subseteq S_2$. 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 $S \subseteq V$, $k_f(S) \leq k_1$ if and only if $k_{f_1}(S) \leq k_1$ and $k_{f_2}(S) \leq k_1$. Thus, for every set $S \subseteq V$, $k_f(S) \leq k_1$ if and only if $S \subseteq S_1$ and $S \subseteq S_2$, or equivalently, if $S \subseteq S_1 \cap S_2$. Hence, the packing dimension of $\mathcal{S} \cap S_2$ is at most $k_1 + k_2$. \square

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

Consider a set $S \subseteq V$ with $|S| > 1$. For the set family $2^V \cap \mathcal{S}^*$, 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 $\text{pdim}(2^V \cap \mathcal{S}^*) \leq 1$ for all $S \subseteq V$ with $|S| > 1$. Note that we have

$$\mathcal{S} = \bigcup_{S \in \mathcal{S}} 2^V \cap \mathcal{S}^*$$

for every downward closed set system \mathcal{S} . Combined with Proposition 132, 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 d sets. We can show that such bounded set families that are downward closed and non-trivial have packing dimension at most $(k \cdot d)^{O(1)}$. Together with the $(\log k)$ hardness [Tre01] of set cover where each set has cardinality at most k and each element appearing in $(\log k)^{O(1)}$ sets, this packing dimension bound gives the hardness $(\log k)^{O(1)}$ for the Vector Bin Packing problem when d is a large constant. Unfortunately, the exponential dependence is necessary for the packing dimension of bounded set systems, and thus, this approach does not yield the optimal $(\log d)$ 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 where each set has cardinality at most k and each element appears in at most d sets, the packing dimension of $\mathcal{S}^\#$ can be upper bounded by $(k \cdot d)^{O(1)}$. Together with the $(\log k)$ hardness of simple k -set cover (proved in Section 8.5), we get the optimal $(\log d)$ 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 8.3.3.

8.3.2 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 at most k and each element appears in at most d sets has packing dimension at most polynomial in k, d .

Theorem 133. 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 ℓ sets. Then,

$$\text{pdim}(\mathcal{S}^\#) \leq (k \ell)^{O(1)}$$

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

We prove the embedding result by writing the set family $\mathcal{S}^\#$ as an intersection of ℓ structured set families each of which has packing dimension at most $O(1)$. We can then upper bound the packing dimension of $\mathcal{S}^\#$ using Proposition 132. The structured set systems we study are sun over-bouquets, which are a disjoint union of sun over-sets that have a single element as the kernel. The formal definition of the sun over-bouquet set families is below. See Figure 8.1 for an illustration.

Definition 134. (Sun over-bouquets) A simple set system on a universe V is called a sun over-bouquet with core $U \subseteq V; U \neq \emptyset$ if the following hold.

1. Every set $S \in \mathcal{S}$ satisfies $|S \setminus U| = 1$. 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 $S_1 \setminus S_2 \neq \emptyset$, we have $S_1 \setminus U = S_2 \setminus U = S_1 \setminus S_2$.

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

Lemma 135. Fix an integer $k \geq 2$. Let \mathcal{S} be a simple set family defined on a universe V that is a sun over-bouquet with core U . Furthermore, each set in the family has cardinality at most k and each element appears in at most ℓ sets. Then, there exists an embedding $\psi : [0, 1]^K$ that satisfies

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

$$|\psi(S)| \leq 1;$$

(B) For every set $S \in \mathcal{S}^\#$ with $S \setminus U \neq \emptyset$,

$$|\psi(S)| > 1;$$

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

$$|\psi(S)| \leq 1;$$

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

$$|\psi(S)| > 1;$$

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

⁶A sun over is a collection of sets S_1, S_2, \dots, S_r whose pairwise intersection is constant i.e., there exists U such that $U = S_i \cap S_j$ for all $i, j \in [r]; i \neq j$. This constant intersection U is called the kernel of the sun over.

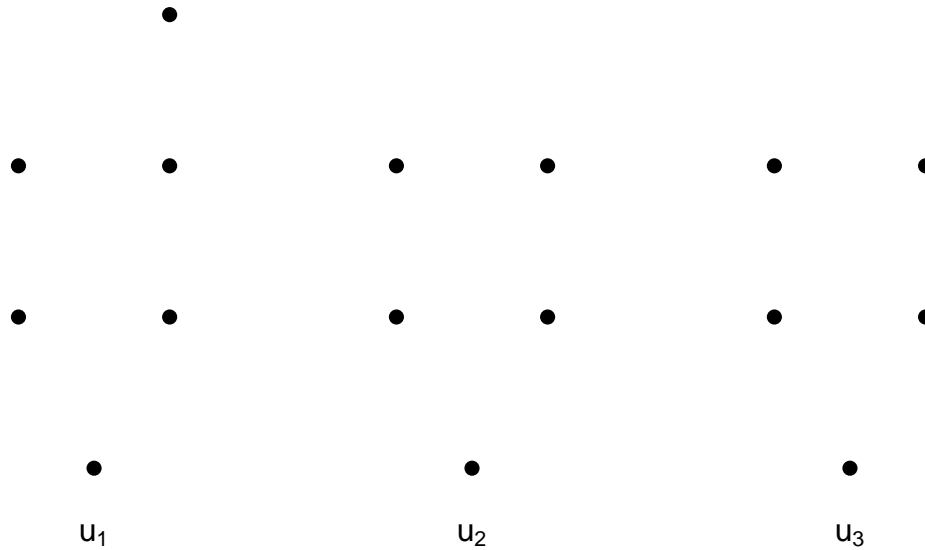


Figure 8.1: An illustration of a sun over-bouquet set family. Here, \mathcal{S} is the family of all the green colored sets. It is a sun over-bouquet with core $U = \{u_1, u_2, u_3\}$. In the embedding, we ensure that the ℓ_1 norm of the left red set is greater than the first step while the right side red set is handled in the second step.

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

$$V_i = \bigcup_{S \ni u_i} S$$

for all $i \in [m]$. Here, \mathcal{S} restricted to $\{u_i\} \cup V_i$ is a sun over set system with a single element u_i as the kernel for every $i \in [m]$. As each set in \mathcal{S} has cardinality at most k and each element appears in at most k sets, we get that $|V_i| \leq k$ for all $i \in [m]$. For every $i \in [m]$, we order the elements of V_i as $v_{i,1}, v_{i,2}, \dots, v_{i,k}$ (with repetitions if needed).

We construct the final embedding f as a concatenation of embeddings with smaller dimensions $f := (f_0; g; g^0)$ where $f_0 : V \rightarrow [0, 1]^2$, $g : V \rightarrow [0, 1]^{K_1}$ and $g^0 : V \rightarrow [0, 1]^{K_2}$ all satisfy the conditions (A) and (C). In other words, for every set $S \in \mathcal{S}$ satisfying either $S \cap U = \emptyset$ and $|S| \leq k$, we have $\|f_0(S)\|_1 \leq 1$, $\|g(S)\|_1 \leq 1$, and $\|g^0(S)\|_1 \leq 1$. Note that for a set $S \subseteq V$, we have

$$\|f(S)\|_1 = \max(\|f_0(S)\|_1; \|g(S)\|_1; \|g^0(S)\|_1)$$

Thus, if f_0, g, g^0 satisfy the conditions (A) and (C), the final embedding f also satisfies the conditions (A) and (C). Furthermore, the parameters K_1 and K_2 are chosen such that the final dimension of f , $2 + K_1 + K_2$ is at most $(k)^{O(1)}$.

First, we define the embedding $f_0 : V \rightarrow [0, 1]^2$ that satisfies the conditions (A) and (C) with the additional property that for every set $S \in \mathcal{S}$ with $|S| > k$, or if $S \cap U \neq \emptyset$ and $S \cap V_0 \neq \emptyset$, or if $|S \cap U| > 1$, we have $\|f_0(S)\|_1 > 1$. We obtain this by a simple two-dimensional embedding

as follows:

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

We can verify that f_0 satisfies the conditions (A) and (C). Furthermore, suppose that $kf_0(S)k_1 > 1$ for a set $S \subseteq V$. Then, we can obtain the following observations that we will use later.

1. As $f_0(v)_1 = 1$ for all $v \in U$, $|S \cap U| \geq 1$. As $f_0(v)_1 = \frac{1}{k}$ for all $v \in V_0$, if $S \cap U = \emptyset$, then $|S \cap V_0| \geq k$.
2. As $f_0(v)_2 = \frac{1}{k}$ for all $v \in V$, $|S| \geq k$. Thus, for every set $S \subseteq V$ such that $|S| > k$, we have $kf_0(S)k_1 > 1$, and hence $kf(S)k_1 > 1$. This already proves the condition (D) of the lemma.

Overview of rest of the proof. We now restrict our attention to sets $S \subseteq V$ such that $|S| \leq k$; $S \cap V_0 = \emptyset$, and $|S \cap U| \geq 1$. Our goal is to find an embedding for these sets that satisfies the conditions (A), (B), and (C). This is the technically challenging part of the proof and requires setting various coordinates carefully to encode the properties of the set system. We break this down into two steps: eliminating the “cross-sun over” sets, and pinning down the “intra-sun over” sets. We give an overview of the ideas used in the two steps before presenting the full proof.

1. The cross-sun over sets are the sets $S \subseteq V$ that contain u_i for some $i \in [m]$, but also intersect another sun over $i \in [m]$, $S \cap V_0 \neq \emptyset$ for an $i \in [m]$. Note that such a set satisfies $S \not\subseteq S^\#$ and $S \cap U \neq \emptyset$, and thus, to satisfy the condition (B), we need to ensure that $kf(S)k_1 > 1$ for such sets. We achieve this by constructing an embedding $g : [0, 1]^{k-1}$ that satisfies the conditions (A) and (C) and has $kg(S)k_1 > 1$ for all cross-sun over sets. We illustrate the idea used in constructing this embedding using a toy example. Suppose that we have a set of pairs of elements $(u_1, v_1); (u_2, v_2); \dots; (u_n, v_n)$, and let their union be denoted by $W = \{u_1, v_1, u_2, v_2, \dots, u_n, v_n\}$. Our goal is to find an embedding $g : W \rightarrow [0, 1]^2$ such that

- (a) $kg(u_i) + g(v_i)k_1 \geq 1$ for all $i \in [n]$.
- (b) $kg(u_i) + g(v_j)k_1 > 1$ for all $i, j \in [n]; i \neq j$.

We construct this embedding by choosing distinct real numbers $\alpha_1, \alpha_2, \dots, \alpha_n \in (0, 1)$ and setting $g(u_i) = (\alpha_i, 1 - \alpha_i)$ and $g(v_i) = (1 - \alpha_i, \alpha_i)$ for all $i \in [n]$. Note that $kg(u_i) + g(v_j)k_1 = 2$ for all i, j , and thus $kg(u_i) + g(v_j)k_1 = 1$ if and only if $g(u_i) + g(v_j) = (1, 1)$, or equivalently, when $i = j$. The actual construction extends this idea with two differences: first, the v_i s could have more than one element, and we use the pairs idea multiple times to account for this, and second, we need to ensure that the sum of the embedding of any k elements in $V \cap (U \cup V_0)$ has ℓ_1 norm at most 1, and thus, we need to choose the embedding g to be $(\alpha_i, 2 - \frac{1}{k} - \alpha_i)$ and set $\alpha_i \in (1 - \frac{1}{k}, 1)$.

2. The intra-sun over sets are the sets $S \subseteq V$ such that $u_i \in S$ and $S \subseteq V_i = \{u_i, v_i\}$ for some $i \in [m]$. We need to ensure that every intra-sun over set S such that $S \not\subseteq S^\#$ satisfies $kf(S)k_1 > 1$. We achieve this by constructing an embedding $g : V \rightarrow [0, 1]^{k-2}$ that

satisfies the conditions (A), (C) and $\text{rk}_1(g(S)) > 1$ for every intra-sun over set S such that $S \not\subseteq S^\#$.

Fix an $i \in [m]$. For every intra-sun over set $S = \{u_i\} \cup V_i$ with $u_i \in S$ and $S \not\subseteq S^\#$, we use a single dimensional embedding $g: V \rightarrow [0, 1]$ that satisfies the conditions (A) and (C) but $\text{rk}_1(g(S)) > 1$. We achieve this by setting $g(u_i) = 1 - \frac{j|S_i| - 1}{k} + \epsilon$, and $g(v) = \frac{1}{k}$ for every $v \in V_i \setminus S$, and $g(v) = 0$ for all $v \in V \setminus S$, where $\epsilon < \frac{1}{k}$ is a small positive constant. Note that $\text{rk}_1(g(T)) = 1$ for every set T such that $T \cap S = \emptyset$, and since $g(v) = \frac{1}{k}$ for all $v \in V \setminus U$, the embedding g satisfies the conditions (A) and (C).

We can construct such a single dimensional embedding for every intra-sun over set S and take their concatenation to obtain the required embedding. However, there could be exponential (in k) number of such intra-sun over sets $S = \{u_i\} \cup V_i$, $u_i \in S$ such that $S \not\subseteq S^\#$. We get around this issue by observing that we need the single dimensional embedding only for minimal intra-sun over sets that don't belong to $S^\#$. In fact, as the set system $S^\#$ restricted to $\{u_i\} \cup V_i$ is a sun over with single element u_i as the kernel, we can deduce that every minimal intra-sun over set S with $S \not\subseteq S^\#$ is of the form $\{u_i, x, y\}$ where $x, y \in V_i$. Thus, we can construct the single dimensional embedding for such sets and take their concatenation to obtain the required embedding with dimension at most $\sum |V_i|^2 (k)^2$.

As every set $S \subseteq V$ such that $S \setminus V_0 = \emptyset$, $|S \setminus U_j| = 1$ is either a cross-sun over set or an intra-sun over set, the two steps together prove that $(f_0; g; g')$ satisfies the conditions (A), (B) and (C).

We now present the full formal proof of the two steps.

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

1. g satisfies the conditions (A) and (C) i.e. for every set $S \subseteq V$, $\text{rk}_1(g(S)) > 1$, and for every set $S \subseteq V$ with $S \setminus U = \emptyset$ and $|S \setminus U_j| = k$; $\text{rk}_1(g(S)) = 1$;
2. For every "cross-sun over" set $S \subseteq V$ with $u_i \in S$ for some $i \in [m]$, and $S \setminus V_{i^0} = \emptyset$ for $i^0 \in [m]$; $i^0 \in i$, we have $\text{rk}_1(g(S)) > 1$.

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

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 $f_l: V \rightarrow [0, 1]^2$, $l \in [k]$ as follows. Consider an $l \in [k]$.

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

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

2. For $i \in [m]$ and $v_{ij} \in V_i$, we set $f_l(v_{ij}) = (0, 0)$ if $v_{ij} \notin V_{i^0}$. We set

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

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 $l \in [k]$.

(A) Consider a set $S \subseteq V$. Let $i \in [m]$ be such that $u_i \in S \setminus U$. We have

$$\begin{aligned} f_l(S) &= \sum_{v \in S} f_l(v) \\ &= \sum_{v \in S \cap U} f_l(v) + \sum_{v \in S \setminus U} f_l(v) \\ &= \sum_{v \in S \cap U} (u_i; 2 \frac{1}{k} - u_i) + \sum_{v \in S \setminus U} (0; 0) \\ &= (1; 1) \end{aligned}$$

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

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

Let $S \subseteq V \setminus V_0$, $\sum_{i \in S} u_i = 1$ be such that

$$\|g(S)\|_1 > 1$$

i.e. $\|f_l(S)\|_1 > 1$ for all $l \in [k]$. Suppose that $S \setminus U = \{u_i\}$. Then, we claim that $\sum_{v \in S \setminus U} f_l(v) = (u_i; 2 \frac{1}{k} - u_i)$. Suppose for contradiction that this is not the case, and there exists $v_0 \in S \setminus U$ with $v_0 \notin \{u_i\}$ and $l \in [k]$ such that $v_0 \in S$. We have

$$\begin{aligned} f_l(S) &= \sum_{v \in S} f_l(v) \\ &= \sum_{v \in S \setminus \{v_0\}} f_l(v) + f_l(v_0) \\ &= (u_i; 2 \frac{1}{k} - u_i) + (0; 0) \\ &= (1 + u_i; 1 + u_i) \end{aligned}$$

As $u_i \notin \{v_0\}$, $\|f_l(S)\|_1 > 1$, a contradiction. Thus, for every set $S \subseteq V$ such that $\sum_{i \in S} u_i = 1$, $S \setminus V_0 \subseteq \{u_i\}$ for some $i \in [m]$, we have $\|g(S)\|_1 \leq 1$.

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

1. g^0 satisfies the conditions (A) and (C).
2. For every $i \in [m]$ and "intra-sun owner" set $S \subseteq \{u_i\} \cup V_i$ such that $\sum_{u \in S} u = 1$ and $S \not\subseteq V_i$, we have $\|g^0(S)\|_1 > 1$.

We achieve this by setting $g^0 = (g_1; g_2; \dots; g_{(k-2)})$ where each $g_l, l \in [(k-2)]$ satisfies the conditions (A) and (C), and the overall function g^0 satisfies the second condition above.

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

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

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

(a) If $V_{i,l} \cap \{u_i\} \neq \emptyset$, we set $g(u_i) = 0$ and $g(v) = 0$ for all $v \in V_i$.

(b) If $V_{i,l} \cap \{u_i\} = \emptyset$, we set $g(v) = \frac{1}{k}$ for all $v \in V_{i,l}$, and $g(v) = 0$ for all $v \in V_i \setminus V_{i,l}$.

We set

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

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

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

(A) Consider a set $S \subseteq V$. Let $u_i \in S \setminus U$. If $u_i \in V_{i,l} \subseteq S^\#$, $g(v) = 0$ for all $v \in S$, and thus we have $g(S) \leq 1$. Now suppose that $u_i \in V_{i,l} \not\subseteq S^\#$. This implies that $V_{i,l}$ is not a subset of S . As $|V_{i,l}| = 2$, $|V_{i,l} \setminus S| \geq 1$. We get

$$\begin{aligned} g(S) &= g(u_i) + \sum_{v \in S \setminus V_i} g(v) \\ &= g(u_i) + \sum_{v \in S \setminus V_i} g(v) \\ &= g(u_i) + \frac{1}{k} \\ &= 1 - \frac{2}{k} + \frac{1}{k^2} + \frac{1}{k} \leq 1 \end{aligned}$$

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

Suppose that a set $S \subseteq V$ satisfies $S \cap V_i \neq \emptyset$ for some $i \in [m]$, and $u_i \in S, S \not\subseteq S^\#$. Then, we claim that $g(S) > 1$. Suppose for the sake of contradiction that $g(S) \leq 1$. Then, we have $g(S) \leq 1$ for all $l \in [(k-2)g]$. 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 $S(v) \subseteq S$ 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 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 $u_i; s_{j_1}; s_{j_2} \notin S^\#$ as otherwise, if there exists $T \subseteq S$ such that $u_i; s_{j_1}; s_{j_2} \in T$, we have $S(s_{j_1}) = S(s_{j_2}) = T$. Let $l \in [(k-2)g]$ be such that $V_{i,l} = \{s_{j_1}; s_{j_2}\}$. As $V_{i,l} \cap \{u_i\} = \emptyset$, we have $g(v) = \frac{1}{k}$ for all $v \in V_{i,l}$ and

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

Thus, we get that

$$\begin{aligned} \sum_{v \in S} g(v) &= g(u_i) + \sum_{v \in S \setminus \{u_i\}} g(v) \\ &= g(u_i) + \sum_{v \in V_{i^0}} g(v) \\ &= 1 + \frac{2}{k} + \frac{1}{k^2} + \frac{2}{k} = 1 + \frac{1}{k^2} \end{aligned}$$

contradicting the fact that $g(S) \leq 1$. 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 $S = S(s_1)$, which implies that $S \in \mathcal{S}^\#$, a contradiction. Thus, for every set $S \subseteq V$ such that $u_i \in S$, $S \cap u_i g[V_i]$ for some $i \in [m]$ and $kg(S)k_1 \leq 1$, we have $S \in \mathcal{S}^\#$.

Final embedding. We define the final embedding $g : V \rightarrow [0, 1]^{2+2k + \binom{k}{2}}$ as $g = (f_0; g; g^0)$. As each of these embeddings satisfies the conditions (A) and (C), the final embedding g also satisfies the conditions (A) and (C).

Suppose that $kg(S)k_1 \leq 1$ for a set $S \subseteq V$. Then, $kg_0(S)k_1 \leq 1$, $kg(S)k_1 \leq 1$ and $kg^0(S)k_1 \leq 1$. Condition (D) follows immediately as $kg_0(S)k_1 \leq 1$ implies that $|S_j| \leq k$.

We now return to condition (B). Suppose that $S \subseteq V$ with $S \setminus U \in \mathcal{E}$; satisfy $kg(S)k_1 \leq 1$. Our goal is to show that $S \in \mathcal{S}^\#$. We have already deduced from $kg_0(S)k_1 \leq 1$ that $|S \setminus U_j| \leq 1$. As $S \setminus U \in \mathcal{E}$, we have $|S \setminus U_j| = 1$, and by using $kg_0(S)k_1 \leq 1$ again, we get that $S \setminus V_0 = \emptyset$. Let $S \setminus U = \{u_i\}$. As $kg(S)k_1 \leq 1$, using the argument in the first step, we can conclude that $S \setminus V_{i^0} = \emptyset$ for all $i^0 \in I$. Thus, $S \cap u_i g[V_i]$. By using the argument in the second step, $kg^0(S)k_1 \leq 1$ implies that $S \in \mathcal{S}^\#$.

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

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

$$\mathcal{T}^\# = \mathcal{S}^\# [f : S \subseteq V \text{ n } U : |S_j| \leq k]$$

Corollary 136. Suppose that \mathcal{T} is a set family defined on a universe V with

$$\mathcal{T} = \mathcal{S} [f : S \subseteq V \text{ n } U : |S_j| \leq k]$$

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

$$\text{pdim}(\mathcal{T}^\#) \leq \binom{k}{2}^{O(1)}$$

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

Proof. As S is a sun over-bouquet, from Lemma 135, there exists an embedding $\iota: [0; 1]^k$ that satisfies the conditions (A); (B); (C) and (D) with $K = (k)^{O(1)}$. Conditions (A) and (C) together imply that

$$k f(S) k_1 \leq 1$$

for all $S \in T$. Note that

$$T^\# = S^\# [f(S) \cup \bigcup_{j \in S} S_j] \text{ kg}$$

Suppose that $S \cap U$ is a subset of V with $|S \cap U| \geq k$. If $S \cap U = \emptyset$, then $|S_j| > k$, which implies that $k f(S) k_1 > 1$ using condition (D). If $S \cap U \neq \emptyset$, then $S \not\subseteq S^\#$ which implies that $k f(S) k_1 > 1$ using condition (B). Thus, $k f(S) k_1 \leq 1$ if and only if $S \in T^\#$. \square

We are now ready to prove our main embedding result i.e. Theorem 133.

Proof of Theorem 133 We define a graph $G = (V; E)$ as follows: two elements $v, v' \in V$ are adjacent in G if there exist sets $S_1, S_2 \in S$ (not necessarily distinct) such that $v \in S_1; v' \in S_2; S_1 \cap S_2 \neq \emptyset$. As the cardinality of each set S is at most k and each element v is present in at most k sets, the maximum degree of a vertex v can be bounded above as

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

Thus, the chromatic number $\chi(G)$ is at most $L = \Delta(G) + 1 = k(k-1)^2 + 1 = k^2 - 2k + 2$. 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 an independent set in G . For every $j \in [L]$, as U_j is an independent set in G , we have

1. For every set $S \in S$, $|S \cap U_j| \leq 1$.
2. Any two sets $S_1, S_2 \in S$ 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 $S_1; S_2; \dots; S_L$ as follows:

$$S_j = \{S \in S : S \cap U_j \neq \emptyset\} \cup \{f(S) \cup \bigcup_{j \in S} S_j\} \text{ kg}$$

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

$$\bigcap_{j \in [L]} S_j = S^\#$$

Using Proposition 132, in order to bound the packing dimension $\rho(S^\#)$, it suffices to bound the packing dimension of $S_j, j \in [L]$.

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

$$S_j = S_j^0 [f(S) \cup \bigcup_{j \in S} S_j] \text{ kg}$$

with

$$S_j^0 = \{S \subseteq U_j : |S| \leq k\}$$

Here, S_j^0 is a simple set system which satisfies the following properties:

1. Each set in S_j^0 has cardinality at most k and each element appears in at most k sets in S_j^0 .
2. Every set $S \subseteq S_j^0$ satisfies $|S \cap U_j| = 1$. As S is non-trivial, for every $u \in U_j$, there exists a set $S \in S_j^0$ with $u \in S$.
3. For every pair of sets $S_1, S_2 \in S_j^0$ 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 S_j^0 is a sunflower bouquet with core U_j . Using Corollary 136, we get that $\dim(S_j^0) = O(k)$ for all $j \in [L]$, which completes the proof. \square

8.3.3 Hardness of Vector Bin Packing

We show that for large enough constant d , Vector Bin Packing is hard to approximate within $(\log d)$. 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 $S_1, S_2, \dots, S_m \subseteq V$ 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. Feige [Fei98] proved a matching hardness of $(1 - 1/e) \ln n$. On set systems where each pair of sets intersect in at most one element i.e. simple instances, the $(\log n)$ 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 137. (Set Cover on simple bounded instances) There exists an irreducible 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 $(\log B)$. Furthermore, in the hard instances, each element occurs in at most $O(B)$ sets.

The details of the parameter modification appear in Section 8.5.

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 120 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 137 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 $O(B)$ sets. We now output a set V of $|V|$ 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 V using m bins.
2. (Soundness.) If there is no set cover of size m in \mathcal{S} , there is no packing of V using m^0 bins.

We use Theorem 133 to compute an embedding $f: V \rightarrow [0, 1]^d$ in polynomial time such that

$$\|f(S)\|_1 \leq 1$$

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

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

Completeness. Suppose that there exist sets $S_1, S_2, \dots, S_m \in \mathcal{S}$ whose union is V . Then, we use 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_1 norm of the sum of the vectors is at most

Soundness. Suppose that the minimum set cover $\mathcal{S}^\#$ has cardinality at least $m^0 + 1$. Then, we claim that the set of vectors V needs $m^0 + 1$ bins to be packed. Suppose for contradiction that there is a vector packing with m^0 bins. In other words, there exists a partition of V into B_1, B_2, \dots, B_{m^0} such that $\|f(B_i)\|_1 \leq 1$ for all $i \in [m^0]$. As $\|f(B_i)\|_1 \leq 1$, $B_i \in \mathcal{S}^\#$ for all $i \in [m^0]$. That is, for every $i \in [m^0]$, there exists a set $S_i \in \mathcal{S}$ such that $B_i \subseteq S_i$. This implies that $\{S_1, S_2, \dots, S_{m^0}\}$ is a set cover of V , a contradiction.

As the original bounded simple set cover problem is hard to approximate within $B = (\log d)$, the resulting Vector Bin Packing is hard to approximate within $(\log d)$. 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

8.4 Vector Scheduling

8.4.1 Monochromatic Clique

In the Monochromatic Clique problem, given a graph $G = ([n]; E)$ and a parameter $k(n)$, the objective is to assign 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 138. (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 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 Coloring problem, which corresponds to the case where $B = 2$. Note that the problem gets easier as B increases. Indeed, when $B > \frac{n}{k}$, we can solve the problem in polynomial time using the canonical SDP relaxation. We present this algorithm and an almost matching integrality gap in Section 8.6.

On the hardness front, we now prove that $\text{MonoChromatic-Clique}(k; B)$ 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 $B = (\log n)$ for some constant > 0 . Next, we amplify this hardness by using lexicographic graph product.

Basic Hardness

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

Lemma 139. For a graph $G = (V; E)$ with $|V| = n$, if $\chi(G) \geq B$, then $\omega(G) \geq n^{1/B}$.

Lemma 140. For a graph $G = (V; E)$ with $|V| = n$, if $\chi(G) \leq B$, then $\omega(G) = O(n^{1/B} \log n)$.

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

Proof. Suppose for contradiction that there is an assignment of k colors to V without a monochromatic clique of size $\geq B$. Using Lemma 140, the subgraphs corresponding to each of the k color classes has chromatic number at most

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

colors. Thus, the whole graph has chromatic number at most $\frac{n}{2^{(\log n)^{1-\epsilon}}}$ colors, a contradiction. \square

Khot [Kho01] proved that assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, the chromatic number of graphs is hard to approximate within a factor of $\frac{n}{2^{(\log n)^{1-\epsilon}}}$ for an absolute constant $\epsilon > 0$. More formally, he proved the following:

Theorem 142. ([Kho01]) There exists a constant $\epsilon > 0$, a function $k = k(n)$, and a randomized reduction that takes as input a SAT instance ϕ on n variables and outputs a graph $G = (V; E)$ with $|V| = N = 2^{\log n^{O(1)}}$ such that

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

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

We observe that Khot's chromatic number hardness immediately gives hardness of Monochromatic Clique.

Lemma 143. There exists a constant $\epsilon > 0$, a function $k = k(n)$ such that the following holds. Assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, given a graph $G = ([n]; E)$, there is a non-deterministic time algorithm for $\text{MonoChromatic-Clique}(k; B)$ when $B = ((\log n)^{1-\epsilon})$.

Proof. Using Khot's reduction, we get that there exists an absolute constant $\epsilon > 0$ such that assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, given a graph $G = ([n]; E)$ and a parameter $k = k(n)$, there is no $n^{(\log n)^{O(1)}}$ time algorithm to distinguish between the following:

1. (Completeness) $\chi(G) \leq k$.
2. (Soundness) $\chi(G) > k \frac{n}{2^{(\log n)^4}}$.

Using Lemma 141, the Soundness condition implies that in any assignment of k colors to G , there is a monochromatic clique of size $\Omega\left(\frac{n}{2^{(\log n)^4}}\right)$, for any constant $\epsilon > 0$. Thus, given a graph G and a parameter k , assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, 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 $\Omega\left(\frac{n}{2^{(\log n)^4}}\right)$.

for any constant $\epsilon > 0$. □

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 144. (Strong Monochromatic-Clique($k; B; C$)) In the Strong Monochromatic-Clique($k; B; C$), given a graph G and parameters $k = k(n); B = B(n); C = C(n)$, 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 colors to the vertices of G , there is a monochromatic clique of size $\Omega(B)$.

We now observe that the chromatic number hardness of Khot [Kho01] implies the same hardness as Lemma 143 for Strong Monochromatic Clique as well.

Lemma 145. There exists a constant $\epsilon > 0$ and a function $k = k(n)$ such that for every constant $C \geq 1$, the following holds. Assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, there is no $n^{(\log n)^{O(1)}}$ time algorithm for Strong Monochromatic-Clique($k; B; C$) when $B = \Omega\left(\frac{n}{2^{C(\log n)^4}}\right)$.

Proof. Note that the function k in Theorem 142 satisfies $k = \Omega\left(\frac{n}{2^{(\log n)^4}}\right)$. Thus, we can replace the soundness condition in Theorem 142 with $\chi(G) > k \frac{n}{2^{C(\log n)^4}}$. Using Lemma 141, this implies that in any assignment of k colors to the vertices of G , there is a monochromatic clique of size $\Omega\left(\frac{n}{2^{C(\log n)^4}}\right)$, where $\epsilon > 0$ is an absolute constant. The hardness of Strong Monochromatic Clique then follows along the same lines as Lemma 143. □

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 146. (Lexicographic product of graphs) Given two graphs G and H , the Lexicographic graph product $G \times H$ 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 $G^2 = G \times G$. 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 $\log_2 n$ times.

Lemma 147. Let $n \geq 2$ be a power of 2. If $\chi(G) = k$, then $\chi(G^n) = k^n$.

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

Lemma 148. 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. Let $f : V(G^2) \rightarrow [k]$ be a given assignment. For a vertex $v \in V(G)$, consider the assignment $f_v : V(G) \rightarrow [k]$ defined as $f_v(u) = f(v; u)$. As every assignment of colors to the vertices of G has a monochromatic clique of size B , there is a color $(v) \in [k]$ and a clique $S(v) \subseteq V(G)$ with $|S(v)| = B$ such that $f_v(u) = (v)$ for all $u \in S(v)$, or in other words $f(v; u) = (v)$ for all $u \in S(v)$. Note that such a set $S(v)$ and (v) exist for $v \in V(G)$. The function $f : V(G^2) \rightarrow [k]$ can also be visualized as an assignment of 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

$$T = \{v \in V(G) : f(v; u) = (v) \text{ for all } u \in S(v)\}$$

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

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

Lemma 149. 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 $G^0 = G^C$. We claim that solving Monochromatic-Clique($k^C; B^C$) on G^0 solves the original Strong Monochromatic Clique problem.

1. (Completeness.) Suppose that $\chi(G) = k$. Then, by Lemma 147, $\chi(G^0) = k^C$.

- (Soundness.) Suppose that in any assignment of colors to the vertices of G , there is a monochromatic clique of size B . Then, by Lemma 148, in any assignment of colors to the vertices of G^0 , there is a monochromatic clique of size B^C . \square

Putting everything together, we obtain the following hardness of Monochromatic Clique. Theorem 150. For every constant $C > 0$, there exists a function $k = k(n)$ such that the following holds. Assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, there is no $n^{(\log n)^{O(1)}}$ time algorithm for Monochromatic-Clique($k; B$) when $B = (\log n)^C$.

Proof. The proof follows directly by combining Lemma 145 and Lemma 149. \square

8.4.2 From Monochromatic Clique to Vector Scheduling

We now prove Theorem 121 using the above hardness of Monochromatic Clique.

Proof of Theorem 121 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 $G = (V = [n]; E)$, parameter 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 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 vectors as the input and the number of machines is equal to k .

We analyze the reduction.

- (Completeness.) Suppose that there exists a proper 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.
- (Soundness.) Suppose that in any assignment of 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

We set $B = (\log n)^C$ for a large constant C to be set later. We choose k from Theorem 150 such that assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, 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 $B \leq (\log d)^{\frac{1}{C+1}}$. Setting $C = \frac{1}{\epsilon} - 1$, we get that d -dimensional Vector Scheduling has no polynomial time ($(\log d)^{\frac{1}{\epsilon}}$) approximation algorithm assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, for every constant $\epsilon > 0$. \square

Remark 151. In [Im+19], Im, Kell, Kulkarni, and Panigrahi also study the r -norm minimization of Vector Scheduling where the objective is to minimize

$$\max_{k \in [d]} \sum_{i=1}^n (L_i^k)^r$$

where L_i^k denotes the load on the machine i in the k th dimension. They gave an algorithm with an approximation ratio $O\left(\frac{\log d}{\log \log d}\right)^{1+\frac{1}{r}}$. Our reduction from Monochromatic Clique gives almost optimal hardness for this variant as well: we get the hardness $(\log d)^{1+\frac{1}{r}}$ assuming $\text{NP}^* \subseteq \text{ZPTIME}(n^{(\log n)^{O(1)}})$, for every constant $\epsilon > 0$.

8.4.3 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 \mathcal{H} and an integer k , the objective is to assign colors to the vertices of \mathcal{H} minimizing the maximum number of monochromatic vertices in an edge. More formally, we study the following decision version of the problem.

Definition 152. (Balanced Hypergraph Coloring.) In the Balanced Hypergraph Coloring problem, given a s -uniform hypergraph \mathcal{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 \mathcal{H} such that in every edge, each color appears at most c times.
2. The hypergraph \mathcal{H} has no proper coloring with k colors i.e., in any assignment of k colors to the vertices of \mathcal{H} , there is an edge all of whose vertices are assigned the same color.

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

Lemma 153. Given a s -uniform hypergraph $\mathcal{H} = (V, E)$ and parameters k, c , there is a polynomial time reduction that outputs a Vector Scheduling instance with 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 \mathcal{H} such that each color appears at most c times in every edge, then there is a scheduling of \mathcal{H} with makespan at most c .
2. (Soundness.) \mathcal{H} has no proper coloring with k colors, then in any scheduling of \mathcal{H} the makespan is at least c .

Proof. Let $d = |E|$. Order the edges of the hypergraph \mathcal{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 to be equal to the number of colors. There is a natural correspondence between the assignment of colors to the vertices of H , $f : V \rightarrow [k]$, and the scheduling where we assign the vector v_i to the machine $\epsilon(i)$. We now analyze our reduction.

1. (Completeness.) If there exists an assignment of colors $f : V \rightarrow [k]$ where each color appears at most t times in each edge, we assign the vector v_i to the machine $\epsilon(i)$. In any dimension $j \in [d]$, at most t 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 t .
2. (Soundness.) If there exists a vector scheduling $\{v_i\}_{i \in V}$ with makespan strictly smaller than t , assign the color $f(i)$ to the i th vertex of the hypergraph. In any edge of the hypergraph, each color appears fewer than t times as the makespan is smaller than t . Thus, $f : V \rightarrow [k]$ is a proper k -coloring of the hypergraph. \square

We prove the hardness results for Vector Scheduling, namely Theorem 122 and Theorem 123 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 the number of edges in the hypergraph H , and the ratio of the makespans in the completeness and soundness is equal to t . Thus, our goal is to prove the hardness of the Balanced Hypergraph Coloring problem where t is as large as possible, as a function of 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 154. Fix an odd prime number $k \geq 3$ and let $\epsilon = \frac{1}{k^8}$. Given a Label Cover instance $G = (V = L \cup R; E; \{L_i, R_i\}_{i \in E})$, there is a polynomial time reduction that outputs a k^2 -uniform hypergraph $H = (V', E')$ with $|V'| = |L| + |R|$ such that

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

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

We defer the proof of Lemma 154 to Section 8.4.4.

Using Lemma 154, 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 $\text{NP}^* \text{DTIME } n^{O(\log \log n)}$ and another NP-hardness result, by using two different hardness results for the Label Cover problem. These two hardness results prove Theorem 123 and Theorem 122 respectively, using Lemma 153.

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

Theorem 155. Assuming $\text{NP}^* \text{DTIME } n^{O(\log \log n)}$, there is no polynomial time algorithm for the following problem. Given a k^2 -uniform hypergraph $H = (V', E')$ with $m = |V'|$ and $k = (\log m)^{(1)}$, distinguish between the following:

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

Proof. By setting $\epsilon = \frac{1}{k^8}$ in Theorem 129, we have a reduction from a SAT problem on variables to the Label Cover problem $G = (V = L \cup R; E; \pi_L; \pi_R)$ with soundness and $|V| \leq n^{O(\log k)}$, $|L| \leq k^{O(1)}$ and $|R| \leq k^{O(1)}$. Using Lemma 154, we can reduce this Label Cover instance to a Balanced Hypergraph Coloring instance (V^0, E^0) with $|V^0| \leq n^{O(\log k)} 2^{k^{O(1)}}$ and $|E^0| \leq n^{O(\log k)} 2^{k^{O(1)}}$. We set $k = (\log n)^{O(1)}$ such that $|V^0| = n^{O(\log \log n)}$ and $|E^0| = n^{O(\log \log n)}$ to obtain the required hardness of Balanced Hypergraph Coloring. \square

The proof of Theorem 123 follows immediately from Theorem 155 and Lemma 153.

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 156. For any constant $\epsilon > 0$, given a k^2 uniform hypergraph $H = (V^0, E^0)$ with $m = |E^0|$ and $k = (\log \log m)^{O(1)}$, it is NP-hard to distinguish between the following:

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

Proof. By setting $\epsilon = \frac{1}{k^8}$ in Theorem 130, we can reduce a SAT instance on variables to a Label Cover instance $G = (V = L \cup R; E; \pi_L; \pi_R)$ with soundness and $|V| \leq n^{1+o(1)} k^{O(1)}$; $|L| \leq 2^{k^{O(1)}}$, $|R| \leq k^{O(1)}$. By using Lemma 154, we can reduce the Label Cover instance to a Balanced Hypergraph Coloring instance (V^0, E^0) with $|V^0| \leq n^{1+o(1)} 2^{k^{O(1)}}$ and $|E^0|$ at most $n^{1+o(1)} 2^{k^{O(1)}}$. We set $k = (\log \log n)^{O(1)}$ to obtain $|V^0| = O(n^2)$; $|E^0| = O(n^2)$.

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

The proof of Theorem 122 follows immediately from Theorem 156 and Lemma 153.

Finally, we remark that if the structured graph version of the Projection Games Conjecture [Mos15] holds, Lemma 154 and Lemma 153 together prove that d -dimensional Vector Scheduling is NP-hard to approximate within a factor of $(\log d)^{O(1)}$.

8.4.4 Proof of Lemma 154

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

Reduction. For ease of notation, let $\ell = |L|$. For every node $v \in L$ of the Label Cover instance, we have a set of ℓ^n vertices denoted by $f_v = f_{v,g} \in [k]^\ell$. The vertex set of the hypergraph is $V^0 = \bigcup_{v \in L} f_v$.

For every $u \in \mathbb{R}$, and k distinct neighbors of u , $v_1, v_2, \dots, v_k \in L$ with projection constraints $\pi_i : [L] \rightarrow [R]; i \in [k]$, consider the set of k^2 vectors x^{ij} for $i \in [k]; j \in [k]$ which satisfy the following: For every $u \in \mathbb{R}$, and for all $i_1, i_2, \dots, i_k \in [k]$ such that $\pi_{i_j}(u) = u$ for all $i \in [k]$, we have

$$f(x^{i_1 i_1}, x^{i_1 i_2}, \dots, x^{i_1 i_k}) = u \quad (8.1)$$

For every such set of k^2 vectors, we add the edge $e = (v_i, x^{ij}) : 1 \leq i, j \leq k$ to E^0 . We can observe that $|V^0| = |L|k^2$ and

$$|E^0| = |L| \sum_{i,j} k^j = |L| \sum_{i,j} k^j = |L|k^2.$$

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

$$c(x) = \sigma(v)$$

We can observe that this coloring satisfies the property that in every edge $e \in E^0$, each color appears at most k times.

Soundness Suppose that there is a proper coloring $c : V^0 \rightarrow [k]$ of the hypergraph \mathcal{H} i.e. in every edge $e = (v_1, v_2, \dots, v_{k^2}) \in \mathcal{H}$, we have

$$|\{i : c(v_i) = c(v_1)\}| > 1$$

Our goal is to prove that there is a labeling to the Label Cover instance that satisfies at least $\frac{1}{k^2}$ fraction of constraints.

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

Lemma 157. (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 x^1, x^2, \dots, x^k that are all assigned the same color such that

$$f(x_i^1, x_i^2, \dots, x_i^k) = [k]$$

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 $x^{v,1}, x^{v,2}, \dots, x^{v,k} \in \mathbb{R}$ such that all these vectors have the same color $c(x^{v,i}) = c(v)$ for all $v \in L; i \in [k]$ for some function $\sigma : L \rightarrow [k]$. Furthermore, there are a set of coordinates $S(v) \subseteq [n]$ with $|S(v)| \geq n - k^3$ such that

$$f(x_i^{v,1}, x_i^{v,2}, \dots, x_i^{v,k}) = [k]$$

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

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

Lemma 158. Let $u \in R$ be a node on the right side of the Label Cover instance. There are a set of labels $S(u) \subseteq R$ such that $|S(u)| \leq k^5$, and for every $v \in L$ that is a neighbor of u with projection constraint $\pi : L \rightarrow R$, we have $S(u) \setminus \pi(S(v)) = \emptyset$.

Proof. Fix a node $u \in R$ on the right side of the Label Cover instance. Let $v_1, \dots, v_\ell \in L$ be the neighbors of u in the Label Cover instance corresponding to the projection constraints $\pi_1, \pi_2, \dots, \pi_\ell$ respectively. As $|S(v_i)| \leq k^3$ for all $i \in [\ell]$, and the constraints π_i are projections, we have $|\pi_i(S(v_i))| \leq k^3$ for all $i \in [\ell]$. Among these subsets $\pi_i(S(v_i))$ of R , let the maximum number of pairwise disjoint subsets be denoted by ℓ_0 . Without loss of generality, we can assume that $S = \{\pi_i(S(v_i)) : i \in [\ell_0]\}$ is a pairwise disjoint family of subsets.

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

$$S(u) = \bigcup_{i \in [\ell_0]} \pi_i(S(v_i))$$

As S is a family of maximum pairwise disjoint subsets, we have $S(u) \setminus \pi_i(S(v_i)) = \emptyset$ for all $i \in [\ell_0]$. Our goal is to bound the size of $S(u)$, which we achieve by bounding ℓ_0 .

We claim that $\ell_0 \leq k(k-1)$. Suppose for contradiction that $\ell_0 > k(k-1)$. This implies that there are $\ell_0 > k(k-1)$ nodes $v_1, v_2, \dots, v_{\ell_0}$ all adjacent to u such that $\pi_i(S(v_i)) : i \in [\ell_0]$ are all pairwise disjoint. Thus, there exists a color $c \in [k]$ and a set of ℓ_0 nodes $w_1, w_2, \dots, w_{\ell_0}$ adjacent to u corresponding to the projection constraints $\pi_1, \pi_2, \dots, \pi_{\ell_0}$ such that $c \in \pi_i(S(v_i))$ for all $i \in [\ell_0]$, and the sets $\pi_i(S(v_i))$ are pairwise disjoint.

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

1. All these vectors are colored the same:

$$c(f(w_i, x^{ij})) = c \quad \forall i, j \in [k]$$

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

$$f(x^{i,1}, x^{i,2}, \dots, x^{i,k}) = [k]$$

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

We claim that these set of vectors satisfy the condition in Equation (8.1). Fix $a \in R$, and $i_1, i_2, \dots, i_k \in L$ such that $a \in \pi_{i_j}(S(v_{i_j}))$ for all $j \in [k]$. As the family of subsets $\pi_i(S(v_i))$ is a pairwise disjoint family, we can infer that there exists at most one $i \in [k]$ such that $i \in S(w_i)$. Note that if $i \in S(w_i)$, then

$$f(x^{ij} : j \in [k]) = [k]:$$

Thus, we have

$$f(i, j) = [k] \quad \forall i, j \in [k]:$$

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

Hence, we can conclude that $\ell_0 \leq k(k-1)$, and thus $|S(u)| \leq k(k-1)k^3 < k^5$. \square

Now, consider the labeling $\sigma : L \rightarrow \Sigma$; where $(v); v \in L$ is chosen uniformly at random from $S(v)$. Similarly, let $\tau : R \rightarrow \Sigma$ is chosen uniformly at random from $S(u); u \in R$. Using Lemma 158, we can infer that for every edge $(v; u)$ in the Label Cover, this labeling satisfies the edge with probability at least $\frac{1}{|S(v)||S(u)|} \geq \frac{1}{k^2}$. By linearity of expectation, this labeling satisfies at least $\frac{1}{k^2}$ fraction of the constraints in expectation. Hence, with positive probability, the labeling satisfies at least $\frac{1}{k^2}$ fraction of the constraints. This concludes the proof of soundness that if \mathcal{H} has a proper coloring, then there exists a labeling σ that satisfies at least $\frac{1}{k^2}$ fraction of the constraints.

8.5 Hardness of simple 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 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 $(\log n)$ 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 $O(\ln \ln B)$ NP-hardness of set cover on instances where each set has cardinality 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 $\frac{1}{32 \cdot 2^{\log^2 B}}$ where ϵ is an absolute constant to be fixed later.

Theorem 159. ([Aro+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 $\sigma : \Sigma_L \rightarrow \Sigma$ that satisfies all the constraints.
2. (Soundness). No labeling σ can satisfy more than $\frac{1}{32 \cdot 2^{\log^2 B}}$ fraction of the constraints.

Furthermore the instances satisfy the following properties:

1. The alphabet sizes $|\Sigma_L|$ and $|\Sigma_R|$ are both upper bounded by $(\log B)^{O(1)}$.
2. The maximum degree 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 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 $d^0 (deg+1)^d$ partitions. Each partition has B^{1-d} parts, where d is a small constant to be fixed later. The partition system is divided into d^0 groups each containing $(deg+1)^d$ partitions. Each group is further organized into $deg+1$ subgroups each of which contains d^0 partitions. Let $P_{g;s;p}$ denote the s th partition in the p th subgroup of the g th group and $C_{g;s;p;k}$ denote the k th set in $P_{g;s;p}$ where $g \in [d^0]$; $s \in [deg+1]$; $p \in [d]$; $k \in [m]$. The partition system satisfies the four properties in Section 4 of [KAR00], the only difference being that the universe now has size B instead of n^0 . Thus, the covering property (Property 4 in [KAR00]) now states that any covering of N with $m \log B$ sets should contain at least $\frac{1}{4}$ 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 settings B is a constant, we just need to show the existence of such a partition system.

We reduce the Label Cover instance in Theorem 159 to a set cover instance. Using 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 SC . The sets in SC , $C_k(v; a)$ are defined exactly as in [KAR00]. The cardinality of each set is at most $B^0 = deg \cdot B = B^2$. Each element is present in at most $ord = O(B)$ sets. The fact that 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^0 in SC . If there is a set cover of size $\frac{1}{2} n^0 m \log B$ in SC , then there is a labeling σ that satisfies a 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 edges as edges having size $m \log B$.

8.6 SDP Relaxation of Monochromatic-Clique

We consider the following SDP relaxation of the graph coloring problem $Color(V; E)$:

$$\begin{aligned} & \text{Minimize } k \\ & \|u_i - u_j\| = 1 \quad \forall (i, j) \in V \\ & \|u_i - u_j\| \leq \frac{1}{k-1} \quad \forall (i, j) \in E \end{aligned}$$

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 \bar{G} . We have the following sandwich property due to [Knu94]:

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

8.6.1 Algorithm when $B > \frac{p}{n}$

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 k in G , and we output YES. If $\chi(G) > k$, 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 k , thus the problem is trivial.

8.6.2 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 k . We now prove that this cannot be significantly improved:

Theorem 160. For n large enough, there exists a graph $G = (V; E)$ with n vertices, and a parameter k such that

1. $\chi(G) \leq k$.
2. In any assignment of k colors to the vertices of G , there is a monochromatic clique of size $\Omega(n^{1/4})$.

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. $\chi(G) \leq k$.
2. In any assignment of k colors to the vertices of the graph G , there exists a monochromatic independent set of size $\Omega(n^{1/4})$.

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

$$\chi(G) \leq 2n^{3/4} + O(n^{1/3} \log n)$$

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

Furthermore, the random graph $G(n; p)$ with $p = \frac{1}{n}$ has $\omega(K_6)$ with high probability (See e.g., [FK15]). Thus, using Lemma 139, we can infer that in any subset of size $\Omega(n^{1/4})$ there is an independent set of size at least $\Omega(n^{1/4})$. Hence, in any assignment of k colors to the vertices of the graph G , there is a monochromatic independent set of size $\Omega(n^{1/4})$. Taking the complement, we get a graph with the required properties. \square

Chapter 9

Approximate hypergraph vertex cover and generalized Tuza's conjecture

9.1 Introduction

The relationship between minimum vertex covers and maximum matchings of graphs and hypergraphs is a fundamental and well-studied topic in combinatorics and optimization. Even though the worst-case factor gap between the two parameters cannot be improved on arbitrary uniform hypergraphs, there are some interesting special cases where the ratio between these quantities is smaller. A classic example of this phenomenon is König's theorem on bipartite graphs, where the sizes of minimum vertex covers and maximum matchings are equal.

For the case of $t = 3$, a notorious open problem capturing this gap on special uniform hypergraphs is Tuza's conjecture [Tuz81; Tuz90], which states that in any graph, the number of edges required to hit all triangles is at most twice the maximum number of edge-disjoint triangles. For a hypergraph \mathcal{H} , let us denote by $\tau(\mathcal{H})$ and $\nu(\mathcal{H})$ the sizes of the minimum vertex cover and maximum matching respectively. Tuza's conjecture is then equivalent to the statement $\tau(\mathcal{H}) \leq 2\nu(\mathcal{H})$ for any 3-uniform hypergraph \mathcal{H} obtained by taking the edges of a graph as its vertices, and the triangles C_3 as its (hyper)-edges. (Taking $G = K_4$ shows that the factor 2 is best possible.) The conjecture has been verified for various classes of graphs such as graphs without $K_{3,3}$ -subdivision [Kri95], graphs with maximum average degree less than 1 [Ful15], graphs with quadratic number of edge disjoint triangles [HR01; Yus12], graphs with treewidth at most 6 [BFG19], and random graphs in the $G_{n,p}$ model [BCD20; KP20]. On general graphs, the current best upper bound on the ratio is a factor of 2.67 due to Haxell [Hax99].

Aharoni and Zerbib [AZ20] introduced an extension of Tuza's conjecture to hypergraphs of larger uniformity. This generalized Tuza's conjecture states that for any t -uniform hypergraph \mathcal{H} , the minimum vertex cover $\tau(\mathcal{H}^0)$ of $\mathcal{H}^0 = \mathcal{H}^{(t-1)}$ is at most $\frac{t+1}{2}$ times that of the maximum matching $\nu(\mathcal{H}^0)$. Here, for a t -uniform hypergraph $\mathcal{H} = (V; E)$, the $(t-1)$ -blown-up hypergraph $\mathcal{H}^0 = \mathcal{H}^{(t-1)}$ is a t -uniform hypergraph whose vertices are the set of all $(t-1)$ sized subsets that are contained in at least one edge of \mathcal{H} and corresponding to every edge $e \in \mathcal{H}$, all the $(t-1)$ -sized subsets e' form an edge in \mathcal{H}^0 . Tuza's conjecture is a special case of their conjecture

when $t = 3$ and H has hyperedges corresponding to the triangles in a graph. As is the case with the original Tuza's conjecture, the conjectured value $\frac{t+1}{2}$ is the best possible gap: when H is the complete t -uniform hypergraph on $(t+1)$ vertices, the $(t-1)$ -blown-up hypergraph $H^0 = H^{(t-1)}$ has $\tau(H^0) = 1$ and $\tau^*(H^0) = \frac{t+1}{2}$.

9.1.1 Fractional Tuza's conjecture and the algorithmic hypergraph Turán problem

A first step towards non-trivially bounding $\tau(H)$ in terms of $\tau^*(H)$ for hypergraphs H from some structured family of hypergraphs is proving a fractional version, i.e., obtaining the same upper bound on the ratio between $\tau(H)$ and $\tau^*(H)$, the fractional maximum matching size. By LP duality, this is equivalent to bounding the ratio between $\tau(H)$ and $\tau^*(H)$, the fractional vertex cover value. As $\tau(H) \tau^*(H) = \tau(H) \tau^*(H)$ for any hypergraph H , establishing the fractional version is a necessary step toward bounding $\tau(H) = \tau^*(H)$. Note that understanding the extremal ratio between τ and τ^* on a given family of hypergraphs is equivalent to bounding the integrality gap of the natural linear programming relaxation of vertex cover on that class of hypergraphs.

Krivelevich [Kri95] proved the fractional version of Tuza's conjecture that $\tau(H^{(2)}) \leq 2 \tau^*(H^{(2)})$ for any 3-uniform hypergraph H . A multi-transversal version of Krivelevich's result is proved in a recent work [Cha+20]. In this work, we prove the fractional version of the generalized Tuza's conjecture (up to $\omega(t)$ factors), establishing a non-trivial upper bound on the LP integrality gap for $(t-1)$ -blown-up hypergraphs.

Theorem 161. For any t -uniform hypergraph H , $\tau(H^0) \leq \frac{t}{2} + 2^{\frac{1}{t-1}} \tau^*(H^0)$, where $\tau(H^0)$ and $\tau^*(H^0)$ are respectively the size of the minimum vertex cover and minimum fractional vertex cover of the blown-up hypergraph $H^0 = H^{(t-1)}$. Furthermore, there is an efficient algorithm to approximate vertex cover of $(t-1)$ -blown-up hypergraphs within $\frac{t}{2} + 2^{\frac{1}{t-1}}$ factor.

The vertex cover problem of $(t-1)$ -blown-up hypergraphs is also intimately connected to the famous Hypergraph Turán Problem [Tur41; Tur61] in extremal combinatorics. In the Hypergraph Turán Problem, the goal is to find the minimum size of a family $\mathcal{F} \subseteq \binom{[n]}{t-1}$ of subsets of $[n]$ with cardinality $(t-1)$ such that for every subset S of $[n]$ of size t , there exists a set $T \in \mathcal{F}$ such that T is a subset of S . The best known upper bound is due to [Sid97]: there exists a family \mathcal{F} of size $O\left(\frac{\log t}{t}\right) \binom{[n]}{t-1}$ such that for every subset S of $[n]$ of size t , there exists a subset $T \in \mathcal{F}$ such that T is contained in S . On the other hand, the lower bound situation is rather dire, with only second-order improvements [CL99; LZ09] over the trivial $\frac{1}{t} \binom{[n]}{t-1}$ lower bound.

Note that the Hypergraph Turán Problem is precisely the minimum vertex cover problem on $H^{(t-1)}$ when H is the complete t -uniform hypergraph. Thus, for a general hypergraph, finding vertex covers on the blown-up hypergraph $H^{(t-1)}$ can be viewed as an algorithmic version of the Hypergraph Turán problem.

Problem 162. (Algorithmic Hypergraph Turán Problem (AHTP)) Given a t -uniform hypergraph $H = (V = [n]; E)$, find the minimum size of a family $\mathcal{F} \subseteq \binom{[n]}{t-1}$ of subsets of V of size $(t-1)$ such that for every hyperedge $e \in E$; there exists $T \in \mathcal{F}$ such that T is a subset of e .




Figure 9.1: The tent

The problem is a generalization of the minimum vertex cover on graphs, which corresponds to the case $t = 2$. As AHTP can be cast as a vertex cover problem on uniform hypergraphs, there is a trivial factor approximation algorithm for this problem. We prove Theorem 161 by obtaining an improved algorithm for AHTP based on rounding the standard LP relaxation of $H^{(t-1)}$.

We now briefly describe this rounding approach. First, using threshold rounding, we argue that one may focus on the case when the LP solution does not have any variables that are assigned values greater than $\frac{1}{2}$. Let S be the set of vertices of H^0 that are assigned non-zero LP value. The thresholding procedure ensures that every hyperedge $e \in E(H^0)$ intersects with S in at least $\frac{t}{2}$ vertices. We can bound the cardinality of S from above by $2 \cdot \text{OPT}$ using the dual matching LP, where OPT is the cost of the optimal LP solution. Our goal then becomes finding a vertex cover of size at most about $\frac{2}{t} \cdot \text{OPT}$. We achieve this by a color-coding technique: we randomly assign a color from $\{0, 1\}^t$ to each vertex of H independently. Most edges of H are almost balanced under this coloring, in the sense that each color appears at least $\frac{1}{t}$ times. We then use this balance property to find a small vertex cover in H^0 .

9.1.2 Vertex cover vs. matching and excluded sub-hypergraphs

The generalized Tuza's conjecture concerns the relationship between $\nu(H)$ and $\tau(H)$ on the $(t-1)$ -blown-up hypergraphs. There have been some works in the literature on the gap between $\nu(H)$ and $\tau(H)$ on other structured class of hypergraphs. An outstanding result of this type is Aharoni's proof [Aha01] that $\nu(H) \leq 2 \cdot \tau(H)$ for all tripartite 3-uniform hypergraphs H . Aharoni and Zerbib [AZ20] asked if there is a structural explanation that unites the generalized Tuza's conjecture and the above result—for example, does the exclusion of a certain substructure in the hypergraph imply better gaps between $\nu(H)$ and $\tau(H)$. A particular substructure they studied is the “tent” subhypergraph (Figure 9.1).

They observed that both tripartite hypergraphs and blown-up 3-uniform hypergraphs cannot contain the tent as a subhypergraph, and asked whether a generalization of Tuza's conjecture might hold for 3-uniform hypergraphs that exclude tents. If this is the case, it could give a common structural explanation of the existence of small vertex covers in uniform hypergraphs.

In this work, we answer this question in the negative. We prove that there are hypergraphs H on n vertices that exclude tents with $\nu(H) \leq (1 - \epsilon)n$. Since $\tau(H) = n/3$ trivially, this shows that the ratio $\frac{\nu(H)}{\tau(H)}$ can approach 3 on tent-free 3-uniform hypergraphs, and the extension of Tuza's conjecture as raised in [AZ20] does not hold. More generally, one might ask if there is some collection of hypergraphs which are excluded from blown-up hypergraphs whose absence

implies a non-trivial gap between α and β . In fact, we prove a stronger statement showing that there is no β -uniform hypergraph family \mathcal{F} (that is absent from blown-up hypergraphs) whose exclusion alone could imply Tuza's conjecture. Our result applies for larger uniformity and the fractional version of Tuza's conjecture.

Theorem 163. For every $\epsilon > 0$ and every finite family of k -uniform hypergraphs \mathcal{F} such that no hypergraph from \mathcal{F} appears in any $(t - 1)$ -blown-up hypergraph $H^0 = H^{(t-1)}$, there is a β -uniform hypergraph T such that T does not contain any hypergraph from \mathcal{F} but $\chi(T) \geq (t - \epsilon) \chi(T)$ (and in fact $\chi(T) \geq (t - \epsilon) \chi(T)$).

The above result rules out the possibility of a “local” proof of Tuza's conjecture. Our construction is a probabilistic one, first sampling each edge of the hypergraph independently with certain probability, and then removing all the copies of hypergraphs in \mathcal{F} using the fact that the family \mathcal{F} satisfies certain sparsity requirements [FM08; BFM10], we can conclude that there is no large independent set in this construction.

We also provide an explicit construction that answers the tent-free question of [AZ20]: our counterexample is the hypergraph with vertex set $[3]^n$ for large enough n and edges being the set of combinatorial lines. By the density Hales Jewett Theorem [FK91; Pol12], there is no large independent set in \mathcal{H} , and using the structure of combinatorial lines, we can prove that \mathcal{H} does not have any tent.

9.1.3 Vertex cover and set cover on simple hypergraphs

As mentioned earlier, AHTP is a special case of vertex cover on uniform hypergraphs. In fact, the blown-up hypergraph $H^{(t-1)}$ is a simple hypergraph: any two edges intersect in at most one vertex. This is simply because any two distinct $(t-1)$ -sized subsets of $[n]$ intersect in at most one vertex. Simple hypergraphs have been well studied in Graph Theory, especially in the context of Erdős-Faber-Rothschild conjecture [Erd81; Erd88] which has been recently proved in a breakthrough result [Kan+21], Ryser's conjecture [Fra+17] and chromatic number of bounded degree hypergraphs [DLR95; FM13].

A natural question is whether we can obtain an approximation ratio smaller than vertex cover on simple hypergraphs. However, Theorem 163 shows that the natural LP has an integrality gap approaching t on simple, and indeed a lot more structured, hypergraphs. But perhaps there are other algorithms that beat the trivial factor t approximation for this problem. We prove that this is not the case, and in fact, vertex cover on simple hypergraphs is as hard as vertex cover on general k -uniform hypergraphs.

Theorem 164. For every $\epsilon > 0$, unless $\text{NP} = \text{BPP}$, no polynomial time algorithm can approximate vertex cover on simple uniform hypergraphs within a factor of $t - \epsilon$. Under the Unique Games conjecture, the inapproximability factor improves to t .

We also study the set cover problem on simple set families where any two sets in the family intersect in at most one element. Equivalently, we want to pick the minimum number of edges to cover all vertices in a simple hypergraph. Kumar, Arya, and Ramesh [KAR00] proved that

¹Simple hypergraphs are also referred to as linear hypergraphs.

set cover problem on simple set systems is hard to approximate within a factor. While this is within constant factor of the approximability of set cover on general set systems, it is natural to wonder if the set cover problem on simple systems is as hard as the same problem on general set systems. Contrary to the vertex cover problem, it turns out that simplicity of the set family does help in getting an improved approximation factor for the set cover—in fact, the greedy algorithm itself delivers such an approximation.

Theorem 165. For set cover on simple set systems over a universe of size n , the greedy algorithm achieves an approximation ratio of $\frac{1}{2} \ln n + 1$. Further, there are simple set systems where the greedy algorithm is off by a factor exceeding $\frac{1}{2} \ln n - 1$.

Interestingly, the dual Maximum Coverage problem, where the goal is to cover as many elements as possible with a specified number of sets, does not become easier on simple set systems and is hard to approximate within a factor exceeding $(1 - e)^{-k}$ [CKL21], the factor achieved by the greedy algorithm on general set systems. In [CKL20], the authors conjecture the hardness of achieving an approximation factor beating $(1 - e)^{-k}$ even for the Maximum Coverage version of AHTP, and call this the Johnson Coverage Hypothesis. They show that this hypothesis implies strong inapproximability results for fundamental clustering problems like k -means and k -median on Euclidean metrics. For example, they showed that the hypothesis implies that k -median is hard to approximate within a factor of 1.73 on ℓ_1 metrics, matching the best hardness factor on general metrics due to Guha and Khuller [GK99].

9.1.4 Other improved hypergraph vertex cover algorithms

Algorithms beating the trivial factor approximation have been obtained for the vertex cover problem on some other families of uniform hypergraphs. In his doctoral thesis, Lovász [Lov75] gave a LP rounding algorithm to obtain a $\frac{1}{2}$ approximation for vertex cover on uniform t -partite hypergraphs. This algorithm is shown to be optimal under the Unique Games Conjecture by Guruswami, Sachdeva, and Saket [GSS15], and an almost matching NP-hardness is also shown. Aharoni, Holzman, and Krivelevich [AHK96] generalized the above algorithmic result to other class of hypergraphs which have a partition of vertices obeying certain size restrictions. A factor $\frac{1}{2}$ approximation algorithm has also been obtained on subdense regular uniform hypergraphs [Car+12].

For the problem of covering all paths of length t (t -Path Transversal), Lee [Lee19] gave a factor $O(\log t)$ approximation. For covering all copies of the star $K_{1,t-1}$, a factor $O(\log t)$ approximation is given in [GL17], and this is tight by a simple reduction from dominating set on degree t graphs. Covering g -connected d -vertex pattern graphs (in particular t -cliques or t -cycles) is as hard as general uniform hypergraph vertex cover [GL17].

9.1.5 Open problems

A number of intriguing questions and directions come to light following our work, and we mention a few of them below.

The most obvious question is whether our algorithm for AHTP can be improved and yield

approximation ratios smaller than 2. Can stronger LP relaxations like Sherali-Adams help in this regard? On the hardness side, essentially nothing is known. There is a straightforward approximation preserving reduction from vertex cover on graphs to AHTP, but this only shows the hardness of beating a factor of 2. Can one show a better inapproximability factor? We do not know any good lower bound on the integrality gap of the LP either—for example, we do not know the existence of a hypergraph for which $(H^{(t-1)}) = (H^{(t-1)})$ grows with t . A natural candidate is the complete t -uniform hypergraph on n vertices for which de Cain conjectured [Cae94] that in fact, $(H^{(t-1)}) = (H^{(t-1)})$ grows with t . However, this is precisely the lower bound of hypergraph Turán problem, and is perhaps very hard to resolve. On the algorithmic side, obtaining $\alpha(\log t)$ approximation algorithm for AHTP would lead to improvements on the hypergraph Turán problem: On the complete hypergraph instance, either the algorithm outputs a family $\mathcal{F}_{t-1}^{[n]}$ of size $\frac{\log t}{t-1} \binom{n}{t-1}$ that covers every subset of size $t-1$, or gives a certificate that any such family should have size at least $\frac{1}{t-1} \binom{n}{t-1}$. In the first case, we get an improvement on the upper bound of hypergraph Turán problem, and in the second case, we resolve de Cain's conjecture.

Similar to the $(t-1)$ -blown-up hypergraphs, one can define the k -blown-up hypergraph of a t -uniform hypergraph—which will be a k^t -uniform hypergraph—and study the vertex cover problem on it. A special case of this problem when $k=2$ is the analog of Tuza's problem for larger cliques, i.e., covering all copies of K_t in a graph by the fewest possible edges. Our algorithm for AHTP extends to this setting, and in particular gives an algorithm with ratio $\alpha(t)$ for the $k=2$ case, beating the trivial $\frac{t}{2}$ factor (see Section 9.3.4 for details). Can one achieve a $\alpha(t^2)$ factor algorithm? A simple reduction from vertex cover on uniform hypergraphs shows an inapproximability factor of $\Omega(1)$, but can one show hardness or integrality gaps?

In general, our work brings to the fore challenges about covering graph structures on both the algorithmic and hardness fronts. On the hardness side, we seem to have essentially no techniques to show strong inapproximability results, as the known PCP techniques where one naturally associates vertices with proof locations do not seem to apply. As mentioned earlier, covering all copies of a t -vertex pattern graph with vertices is as hard to approximate as general t -uniform hypergraph vertex cover which is 2-connected [GL17].

Our structural results show that the LP integrality gap (and therefore also the vertex cover to matching ratio) remains close to 1 on hypergraphs that exclude subgraphs absent $(t-1)$ -blown-up hypergraphs, and thus no “local” proof of Tuza-type conjectures is possible. Are there interesting families of uniform hypergraphs such that vertex cover admits non-trivial approximation (with ratio less than $\frac{t}{t-1}$) on F -free t -uniform hypergraphs?

For the maximization version of AHTP, where we seek to pick a specified number of $(t-1)$ -sized subsets to cover the largest number of edges in a uniform hypergraph, is there an algorithm that beats the $(1 - \frac{1}{t})$ factor (achieved by greedy for the general Max Coverage problem)? The Johnson Coverage Hypothesis of [CKL20] asserts that for any $\epsilon > 0$, a $(1 - \frac{1}{t} + \epsilon)$ -approximation is hard to obtain for large enough t compared to $\frac{1}{\epsilon}$.

We have considered covering problems in this work, and there are interesting questions concerning the dual packing problems as well. For instance, what is the approximability of packing edge-disjoint copies of say K_t -cliques, in a graph? This is a special case of the matching

problem on 2-blown-up hypergraphs. For the maximum matching problem on general uniform hypergraphs, also known as set packing, Cygan [Cyg13] gave a local search algorithm that achieves an approximation factor of $\frac{k+1}{3}$ for any constant $k > 0$. Can we get better algorithms for the maximum matching problem on blown-up hypergraphs?

On the hardness front, k -set packing is inapproximable to a $(k - \log k)$ factor [HSS06]. Known inapproximability results for the independent set problem on graphs with maximum degree k [AKS11; Cha16] imply that the maximum matching problem on uniform simple hypergraphs is hard to approximate within a $\frac{k}{\log^2 k}$ factor. Could maximum matching on simple hypergraphs be easier to approximate than general hypergraphs?

Organization. In Section 9.2, we introduce some notation and definitions. In Section 9.3, we describe and analyze our algorithm for AHTP and prove Theorem 161. Then, in Section 9.4, we prove that the analog of (generalized) Tuza's conjecture does not hold based only on local forbidden sub-hypergraph characterizations, proving Theorem 163, and also giving an explicit construction for the tent-free case posed in [AZ20]. Finally, in Section 9.5, we consider simple hypergraphs and prove Theorems 164 and 165.

9.2 Preliminaries

Notation. We use $[n]$ to denote the set $\{1, 2, \dots, n\}$. We use \mathbb{Z}_n to denote the set $\{0, 1, \dots, n-1\}$. For a set S and an integer $k \leq |S|$, we use $\binom{S}{k}$ to denote the family of all k -sized subsets of S . A hypergraph $H^0 = (V^0, E^0)$ is called a subhypergraph of $H = (V, E)$ if $V^0 \subseteq V$ and $E^0 \subseteq E$. For a hypergraph $H = (V, E)$, we use $\tau(H)$ and $\nu(H)$ to denote the size of the minimum vertex cover and the maximum matching respectively. Similarly, we use $\tau^*(H)$ to denote the minimum fractional vertex cover of

$$\tau^*(H) = \min_{x_v : x_v \in \mathbb{R}_+, \sum_{v \in V} x_v = 1} \sum_{e \in E} \max_{v \in e} x_v$$

We define the k -blown up hypergraph formally:

Definition 166. For a t -uniform hypergraph $G = (V, E)$ and for an integer $1 \leq k < t$, we define the k -blown up hypergraph $H = G^{(k)} = (V^0, E^0)$ as follows:

1. The vertex set $V^0 = \binom{V}{k}$ is the set of all k -sized subsets of V that are contained in an edge of G :

$$V^0 = \{U : U \subseteq V; |U| = k; \exists e \in E : U \subseteq e\}$$

2. For every edge $e \in E$, we include in E^0 all the k -sized subsets of e so that

$$E^0 = \{e^0 : e^0 \subseteq e, |e^0| = k; e \in E\}$$

We will need the following Chernoff bound:

Lemma 167. (Multiplicative Chernoff bound) Suppose X_1, X_2, \dots, X_n are independent random variables taking values in $[0, 1]$. Let $X = X_1 + X_2 + \dots + X_n$, and let $\mu = E[X]$. Then, for any $0 < \epsilon < 1$,

$$\Pr(X \geq (1 + \epsilon)\mu) \leq e^{-\frac{\epsilon^2 \mu}{2}}$$

9.3 LP rounding algorithm for AHTP

In this section, we present our algorithm for AHTP and prove Theorem 161. Given a uniform hypergraph G as an input to the AHTP, let $H = G^{(t-1)}$ be the $(t-1)$ -blown-up hypergraph of G .

9.3.1 Color-coding based small vertex cover

We first prove a lemma that in an $(t-1)$ -blown-up hypergraph $H = ([n]; E)$, there is a vertex cover of size at most $O\left(\frac{\log t}{t}\right)n$ using a color-coding argument. This lemma illustrates the color-coding idea well, and is also useful later in the context of structural characterization of the blown-up hypergraphs. This lemma is not used in the main algorithm, and the reader can skip to Section 9.3.2 for the algorithm.

Lemma 168. Suppose $G = ([n]; E(G))$ is a t -uniform hypergraph and $H = G^{(t-1)} = (V(H); E(H))$. Then, there exists a randomized polynomial time algorithm that outputs a vertex cover with expected size at most $|V(H)| \frac{2 \ln t}{t} + O\left(\frac{1}{t}\right)$.

Proof. Our algorithm is based on the color-coding technique used to get upper bounds for the hypergraph t -uniform problem [KR83; Sid95]. Let $P = \frac{t-1}{2 \ln t}$. Color each vertex of G with $c: [n] \rightarrow [P]$ uniformly independently at random. For $v \in V(H)$ and $i \in [P]$, let $C_i(v)$ denote the number of nodes of v that are colored with i , i.e., $C_i(v) := |\{j \in v : c(j) = i\}|$.

We define a function $f: V(H) \rightarrow \mathbb{Z}_P$ as

$$f(v) = C_1(v) + 2C_2(v) + \dots + (P-1)C_{(P-1)}(v) \pmod{P}$$

For an element $i \in \mathbb{Z}_P$, let $f^{-1}(i)$ denote the set $\{v \in V(H) : f(v) = i\}$. Let $p \in \mathbb{Z}_P$ be such that $|f^{-1}(p)| \geq |f^{-1}(i)|$ for all $i \in \mathbb{Z}_P$. Note that by definition $|f^{-1}(p)| \geq \frac{|V(H)|}{P}$. Let $U \subseteq V(H)$ be defined as follows:

$$U = \{v \in V(H) : \exists i \in [P] \text{ such that } C_i(v) = 0\}$$

We claim that $S = f^{-1}(p) \cup U$ is a vertex cover of H . Consider an arbitrary edge $e = \{v_1, v_2, \dots, v_t\} \in E(H)$. Let the corresponding edge $g \in E(G)$ be equal to $e(G) = \{v_j \in V_j : j \in [t]\}$ where u_1, u_2, \dots, u_t are elements of $[n]$. Without loss of generality, let $v_j = e(G) \cap u_j$. For a color $i \in [P]$, let $C_i(e) = |\{j \in [t] : c(u_j) = i\}|$. We consider two cases separately:

1. First, if there exists a color $i \in [P]$ such that $C_i(e) = 0$, then for every $j \in [t]$, $C_i(v_j) = 0$, and thus, for every $j \in [t]$, $v_j \in U$, and thus $e \subseteq S$.

2. Suppose that for every color $c \in [P]$, $C_i(c) > 0$. We define $f(c) \in \mathbb{Z}_P$ as

$$f(c) = C_1(c) + 2C_2(c) + \dots + (P-1)C_{(P-1)}(c) \pmod P$$

Note that for every $j \in [t]$, we have

$$f(v_j) = f(c(u_j)) \pmod P$$

As the size of $\{c(u_1); c(u_2); \dots; c(u_t)\}$ is equal to P , the size of the set $\{f(v_1); f(v_2); \dots; f(v_t)\}$ is equal to P as well. Thus, there exists $j \in [t]$ such that $f(v_j) = p$ which implies that $v_j \in S$.

Thus, our goal is to upper bound the expected value of $|S|$. Note that $P \leq \frac{t-1}{\ln t}$. By taking union bound over all the colors, we get

$$E[|U|] \leq P \leq \frac{1}{P} \sum_{j=1}^t |V(H)_j| \leq \frac{t-1}{\ln t} e^{2 \ln t} |V(H)| = \frac{1}{\ln t} |V(H)| + O\left(\frac{1}{t} |V(H)|\right)$$

Thus, the expected value of $|S|$ is at most $\frac{1}{\ln t} |V(H)| + E[|U|]$ which is at most $\frac{2 \ln t}{t-1} + O\left(\frac{1}{t} |V(H)|\right)$: \square

9.3.2 LP rounding based algorithm for AHTP

Consider the standard LP relaxation for vertex cover in

$$\begin{aligned} & \text{Minimize} && \sum_{v \in V(H)} x_v \\ & \text{such that} && \sum_{v \in E(H)} x_v \leq 1 \quad \text{and} \quad x_v \in [0, 1] \quad \forall v \in V(H) \end{aligned}$$

Let \bar{x} be an optimal solution to the above Linear Program, and $\text{OPT} = \sum_{v \in V(H)} \bar{x}_v$. Let $S \subseteq V(H)$ be the set of vertices that are assigned positive LP value i.e.

$$S = \{v \in V(H) : \bar{x}_v > 0\}$$

We need a lemma relating $|S|$ and OPT :

Lemma 169. The cardinality of S is at most OPT .

Proof. Consider the dual of the vertex cover LP:

$$\begin{aligned} & \text{Maximize} && \sum_{e \in E(H)} y(e) \\ & \text{such that} && \sum_{e \in V(H)} y(e) \leq 1 \quad \text{and} \quad y(e) \in [0, 1] \quad \forall e \in E(H) \end{aligned}$$

Let \bar{y} be an optimal solution to the above matching LP. By LP-duality, we get $\sum_{e \in E(H)} \bar{y}_e = \text{OPT}$. Recall that for all $v \in V$, $\sum_{e \in E(H)} \bar{y}_e = 1$. By the complementary slackness conditions, we get that for all $v \in V$, $\sum_{e \in E(H)} \bar{y}_e = 1$. Summing over all $v \in V$, we obtain

$$\sum_{v \in V} \sum_{e \in E(H)} \bar{y}_e = \sum_{v \in V} 1 = |V| = \sum_{e \in E(H)} \bar{y}_e = \text{OPT}$$

□

In general, OPT could be much smaller than $|V(H)|$, and thus we cannot use Lemma 168 directly to obtain algorithm for AHTP. However, we can obtain a simple $(t-1)$ -factor approximation algorithm for AHTP using Lemma 168, extending the proof of fractional Tuza's conjecture of Krivelevich [Kri95]. We consider two different cases:

1. Suppose that there is a vertex $v \in V(H)$ such that $\bar{x}_v = 0$. Consider an arbitrary edge $e \in E(H)$ with $v \in e$. As $\sum_{u \in e} \bar{x}_u = 1$, we can infer that there is a vertex $u \in e$ such that $\bar{x}_u \geq \frac{1}{t-1}$. We round \bar{x} to 1 i.e. add v to our vertex cover solution, delete all the edges containing v and recursively proceed.
2. Suppose that for every vertex $v \in V(H)$, we have $\bar{x}_v > 0$. In this case, using Lemma 168, we can find a vertex cover of size $O(\frac{\log t}{t} |V(H)|)$, which can be bounded above by $O(\log t) \text{OPT}$ using Lemma 169.

We now describe a randomized algorithm to round the LP to obtain an integral solution whose expected size is at most $\frac{t}{2} + 2 \sqrt{t \ln t} \cdot \text{OPT}$. As is evident from the second case in the above $(t-1)$ -factor algorithm, the problem is easy when the set of vertices $V(H)$ with non-zero LP value is large. Instead of considering the two different cases based on whether $V(H)$ or not, we take a more direct approach by finding a vertex cover of size $O(1) |V(H)|$. Combined with Lemma 169, we get our required approximation guarantee.

For ease of notation, let $\epsilon = \frac{t}{2} + 2 \sqrt{t \ln t}$. Our first step is to round all the variables above a certain threshold ϵ (Algorithm 1). However, we need to do it recursively to ensure that we can bound the optimal value of the remaining instance.

Algorithm 1 Recursive thresholding for AHTP

- 1: Let $\epsilon = \frac{t}{2} + 2 \sqrt{t \ln t}$.
 - 2: Let \bar{x} be an optimal solution of the LP and let $V^0 = \{v \in V : \bar{x}_v \geq \epsilon\}$.
 - 3: Let $U = V^0$.
 - 4: while V^0 is non-empty do
 - 5: Delete V^0 from $V(H)$, and delete all the edges $e \in E(H)$ that contain at least one vertex $v \in V^0$.
 - 6: Solve the LP with updated H . Update \bar{x} to be the new LP solution.
 - 7: Update $V^0 = \{v \in V(H) : \bar{x}_v \geq \epsilon\}$. Update $U = U \cup V^0$.
 - 8: Output U and the updated H .
-

Let the final updated hypergraph H when Algorithm 1 terminates be denoted by H^0 . Let the optimal cost of the solution \bar{x} for the vertex cover on H^0 be denoted by OPT^0 . We prove that the size of the vertex cover output by the algorithm is not too large:

Lemma 170. When the above recursive thresholding algorithm (Algorithm 1) terminates, we have $|U_j| \leq t^0 \cdot OPT^0 / OPT_{new}$.

Proof. We will inductively prove the following: after line 6 in the while loop of the algorithm, $|U_j| \leq t^0 \cdot (OPT_{old} - OPT_{new})$ where OPT_{new} is the cost of the current optimal solution \bar{x} . Let \bar{x}^0 be the optimal solution before deleting V^0 from H . Let OPT_{old} be the cost of the solution \bar{x}^0 . By inductive hypothesis, we have $|U_j| \leq t^0 \cdot (OPT_{old} - OPT_{new})$.

We claim that $|U_j| \leq t^0 \cdot (OPT_{old} - OPT_{new})$. As \bar{x} is an optimal vertex cover of H , we have that \bar{x}^0 restricted to H has cost at least OPT_{new} . This implies that $\sum_{v \in V^0} \bar{x}_v^0 \geq OPT_{old} - OPT_{new}$. As each $\bar{x}_v^0; v \in V^0$ is at least $\frac{1}{t^0}$, we obtain the required claim. \square

We are now ready to state our main algorithm for APX . The input to the algorithm is a t -uniform hypergraph G , and the output is a vertex cover for the hypergraph $G^{(t-1)}$.

Algorithm 2 Main algorithm

- 1: Apply Algorithm 1 to obtain U and let $H^0 = (V(H^0); E(H^0))$ be the updated H . Let \bar{x} be an optimal solution of the vertex cover LP on H^0 with $\bar{x}_v = \bar{x}_v$ for all $v \in V(H^0)$.
- 2: Let $S = \bigcup V(H^0)$ be defined as $S = \{v \in V(H^0) : \bar{x}_v > 0\}$.
- 3: Let $\epsilon = \frac{4 \ln t}{t-1}$.
- 4: Color the vertices $\{v\}$ of G using $c: [n] \rightarrow \{0, 1\}$ uniformly and independently at random.
- 5: For a vertex $v \in S$ and a color $i \in \{0, 1\}$, let $C_i(v)$ denote the number of nodes that are colored with the color i , i.e. $C_i(v) = |\{j \in S : c(j) = i\}|$. Recall that $|S| = \binom{[n]}{t-1}$.

$$C_i(v) = |\{j \in S : c(j) = i\}|$$

- 6: Let $S^0 \subseteq S$ be defined as the set of vertices v where the discrepancy between two colors is high:

$$S^0 = \{v \in S : |C_0(v) - C_1(v)| \geq (1 - \epsilon) \frac{t-1}{2}\}$$

- 7: We now define a function $f: S \rightarrow \{0, 1\}$ as $f(v) = C_1(v) \bmod 2$.
 - 8: For $i \in \{0, 1\}$, let $f^{-1}(i)$ denote the set of all the vertices $v \in S$ such that $f(v) = i$.
 - 9: Let $p \in \{0, 1\}$ be such that $|f^{-1}(p)| \geq \frac{1}{2} |S^0|$.
 - 10: Let $T \subseteq S$ be defined as $T = S^0 \cap f^{-1}(p)$.
 - 11: Output $T \cup U$.
-

9.3.3 Analysis of the algorithm and proof of Theorem 161

We will first prove that Algorithm 2 indeed outputs a valid vertex cover H .

Lemma 171. $T \cap U$ is a vertex cover of H .

Proof. It suffices to prove that T is a vertex cover of H^0 .

Consider an arbitrary edge $e = (v_1; v_2; \dots; v_t) \in E(H^0)$ corresponding to the edge $e(G) = (u_1; u_2; \dots; u_t) \in E(G)$. Since $v_j \in T$ for all $v \in V(H^0)$, we can deduce that $|e \cap S| \geq t - 1 = t^0$.

Our goal is to show that there exists $s \in [t]$ such that $v_s \in T$. We consider two separate cases:

Case 1: If there is a color $i \in \{0, 1\}$ such that there are at most $(1 - \frac{1}{2})^{t-1}$ nodes of color i in $e(G)$, then for all $j \in [t]$, $C_i(v_j) \leq (1 - \frac{1}{2})^{t-1}$. Since $e \cap S$ is non-empty, there exists $s \in [t]$ such that $v_s \in S$. By definition of S^0 , this implies that $v_j \in S^0$ as well, and thus $S \cap T \neq \emptyset$.

Case 2: Suppose that in the coloring e , both the colors $0, 1$ occur at least $(1 - \frac{1}{2})^{t-1}$ times in e . Let $e^0 = e \cap S$ and let $k = |e^0| - t^0$. Without loss of generality, let $e^0 = (v_1; v_2; \dots; v_k) \in E(H^0)$. For every $j \in [k]$, let $v_j = e(G) \cap (u_1; \dots; u_j)$ for $u_j \in [n]$. First, we claim that $k < (1 - \frac{1}{2})^{t-1}$. We have

$$\begin{aligned} t - k &\geq (1 - \frac{1}{2})^{t-1} - t^0 = (1 - \frac{1}{2})^{t-1} - (1 - \frac{1}{2})^t = \frac{1}{2} (1 - \frac{1}{2})^{t-1} \\ &= \frac{1}{2} (1 - \frac{1}{2})^{t-1} < 0 \end{aligned}$$

Since each color occurs at least $(1 - \frac{1}{2})^{t-1}$ times in $e(G)$, using the above, we can infer that

$$|e \cap S| \geq t - k \geq t - (1 - \frac{1}{2})^{t-1} > t - 1 = t^0$$

We define the value $f(v)$ in the same fashion as we have defined $f(e)$ for $v \in S$: For $i \in \{0, 1\}$, let $C_i(e)$ denote the number of nodes $v \in [t]$ such that $C_i(v) = i$, and let $f(e) = C_1(e) \pmod 2$. Using this definition, we get

$$f(v_j) = f(e) - C_1(v_j) \pmod 2 \quad j \in [k]:$$

As $(u_1; u_2; \dots; u_k) \in E(G)$, we have $f(u_1); f(u_2); \dots; f(u_k) \in \{0, 1\}$ as well. Thus, there exists $s \in [k]$ such that $f(u_s) = 1$, which proves that $v_s \in T$. \square

Note that the expected number of nodes of each color $0, 1$ in a vertex $v = (u_1; u_2; \dots; u_{t-1}) \in S$ is equal to $\frac{t-1}{2}$. The set S^0 is the set of vertices $v \in S$ where there is a color that occurs much fewer than its expected value. We prove that this happens with low probability:

Lemma 172. The expected cardinality of S^0 is at most $\frac{2}{t} |S|$.

Proof. Let $v = (u_1; u_2; \dots; u_{t-1}) \in S$ be an arbitrary vertex in S , where $u_1; u_2; \dots; u_{t-1}$ are elements of $[n]$. For a color $i \in \{0, 1\}$, let the random variable $X(i)$ denote the number of nodes $j \in [t-1]$ such that $C_i(u_j) = i$. We can write $X(i) = \sum_{j \in [t-1]} X(i; j)$, where $X(i; j)$ is the indicator random variable of the event that $C_i(u_j) = i$. We have $\mu = E[X(i)] = \frac{t-1}{2}$. Using multiplicative Chernoff bound (Lemma 167), we can upper bound the probability that $X(i) \leq (1 - \frac{1}{2}) \frac{t-1}{2}$ by

$$\Pr[X(i) \leq (1 - \frac{1}{2}) \frac{t-1}{2}] \leq e^{-\frac{2(t-1)}{4}}$$

For the choice $p = \frac{4 \ln t}{t-1}$, the above probability is at most $\frac{1}{t}$. By applying union bound over the two colors and adding the expectation over all the vertices S , we obtain the lemma. \square

Finally, we bound the expected size of the output of the algorithm:
 Lemma 173. The expected cardinality $E|T \cap U|$ is at most $\frac{t}{2} + 2 \sqrt{\frac{p}{t \ln t}} \text{OPT}$.

Proof. Note that by definition, $\sum_j |S_j| = \frac{1}{2} \sum_j |S_j|$. We bound the expected size of the output of the algorithm $T \cap U$ as

$$\begin{aligned} E|T \cap U| &= E|T| + E|U| - E|S^c| + \frac{1}{2} \sum_j |S_j| + E|U| \\ &= \frac{1}{2} + \frac{2}{t} \sum_j |S_j| + E|U| \quad (\text{Using Lemma 172}) \\ &= \frac{t}{2} + 2 \text{OPT}' + E|U| \quad (\text{Using Lemma 169}) \\ &= \frac{t}{2} + 2 \sqrt{\frac{p}{t \ln t}} \text{OPT} \quad (\text{Using Lemma 170}): \end{aligned}$$

\square

Lemma 171 and Lemma 173 together imply Theorem 161.

9.3.4 $(t; 2)$ -version of AHTP

An interesting generalization of AHTP is the $(t; k)$ -version, the problem of vertex cover on the k -blown-up hypergraph $H = G^{(k)}$ for a t -uniform hypergraph G , for an arbitrary $1 \leq k < t$. The case of $k = 1$ is the standard vertex cover on uniform hypergraphs, and $k = t - 1$ is the AHTP. Note that there is a trivial $\frac{t}{k}$ -factor approximation algorithm for this problem as it can be cast as an instance of vertex cover on $\frac{t}{k}$ -uniform hypergraph. The above algorithm can be shown to achieve a $\frac{t}{k} \alpha(k)$ approximation guarantee for the general problem where $\alpha(k) = \frac{1}{2} + o(1)$ as $k \rightarrow t - 1$.

We now turn our attention to the interesting case of $k = 2$. When the hypergraph G consists of t -cliques in a graph, the vertex cover problem $G^{(2)}$ is the generalization of Tuza's problem where we try to hit all t -cliques with the fewest possible edges. Note that in this case, the trivial hypergraph vertex cover algorithm achieves $\frac{t}{2}$ -factor approximation. We describe how a simplified version of our algorithm can be used to get $\frac{t}{2}$ -factor guarantee: Let $H = G^{(2)} = (V(H); E(H))$, and we iteratively solve the Vertex Cover LP on H to round all the vertices with value at least $\frac{1}{2}$. In the remaining instance, we let $S = V(H)$ to be the vertices of H that are assigned non-zero LP value i.e. $S = \{v \in V(H) : x_v > 0\}$. We use a color coding function $c: [n] \rightarrow \{0, 1\}$ picked uniformly and independently at random, and we output all the vertices $T = \{i, j \in S : c(i) = c(j)\}$. The expected size $E|T|$ is at most $\frac{1}{2} \sum_{i,j \in S} 1 = \frac{1}{2} |S|^2$ OPT as the cardinality of S is at most $\frac{t}{2}$ OPT.

We now argue that Γ is indeed a vertex cover of H . Consider an edge $e = \{v_{i,j} : i \in j \subseteq [t]\} \in E(H)$ corresponding to the edge $e = (u_1; u_2; \dots; u_t) \in E(G)$. Recall that every element of e corresponds to a subset of size $t/2$ of e , and thus, without loss of generality, let $e = \{u_i; u_j\} \in e$ for all $i, j \subseteq [t]$. As $x_{v_{i,j}} < \frac{4}{t^2}$ for all $i, j \subseteq [t]$, there are greater than $\frac{t^2}{4}$ pairs of indices $i, j \subseteq [t]$ such that $x_{v_{i,j}} > 0$, or equivalently $v_{i,j} \in S$. Thus, $|e \cap S| > \frac{t^2}{4}$. For every function $c : [t] \rightarrow \{0, 1\}$, the number of pairs of indices $i, j \subseteq [t]$ such that $c(i) \neq c(j)$ is at most $\frac{t^2}{4}$. Thus, there are at least $\frac{t^2}{2} - \frac{t^2}{4} = \frac{t^2}{4}$ pairs of indices $i, j \subseteq [t]$ such that $c(u_i) = c(u_j)$. As $|e \cap S| > \frac{t^2}{4}$, there exists a pair of indices $i, j \subseteq [t]$ such that $v_{i,j} \in S$, $c(u_i) = c(u_j)$, which implies that $v_{i,j} \in T$. Thus, for every edge $e \in E(H)$ of H , there exists an element $v \in T$ such that $v \in e$, which completes the proof that Γ is a vertex cover of H .

As a corollary of the algorithm, we deduce that for all hypergraphs $G^{(2)}$ of a t -uniform hypergraph G ,

$$(H) \leq \frac{t^2}{4} (H)$$

This proves the fractional version of a conjecture due to Aharoni and Zerbib [AZ20] (Conjecture 1.4) for the case when $k = 2$.

9.4 Forbidden sub-hypergraphs and Tuza's conjecture

Since AHTP is the problem of vertex cover of $H = G^{(t-1)}$ for a given t -uniform hypergraph G , an interesting question is to characterize the uniform hypergraphs H that can arise as the blown-up hypergraph $G^{(t-1)}$ of some t -uniform hypergraph G . A very simple necessary condition is that the hypergraph H should be simple. However, this is not sufficient—there are simple t -uniform hypergraphs H that cannot be written as $H = G^{(t-1)}$ for any G . For example, the t -tent hypergraph (Definition 174) is a simple hypergraph, but cannot be written as a $(t-1)$ -blown-up hypergraph. A natural question in this context is the following:

Is there a finite set of hypergraphs \mathcal{F} such that every hypergraph that does not have any member of \mathcal{F} as a sub-hypergraph can be represented as $G^{(t-1)}$ for some t -uniform hypergraph G ?

In addition to its inherent structural interest, the above question can shed light on Tuza's conjecture. Recall that Aharoni and Zerbib [AZ20] proposed a generalization of Tuza's conjecture stating that $G^{(2)} \leq 2 G^{(2)}$ for all 3-uniform hypergraphs G . They suggested that understanding the structure of blown-up hypergraphs, and specifically, the sub-hypergraphs that it excludes might be a promising approach to establish this conjecture. In particular, they observed that the blown-up hypergraphs do not contain “tents” as a sub-hypergraph.

Definition 174. A t -tent (Figure 9.1) is a set of four t -uniform edges $e_1; e_2; e_3; e_4$ such that

1. $\bigcap_{i=1}^3 e_i \in \mathcal{E}$.
2. $|e_4 \cap e_j| = 1$ for all $i \in [3]$.
3. $e_4 \cap e_i \in e_4 \cap e_j$ for all $i \in j \subseteq [3]$.

In [AZ20], the authors pose the following question. Note that an answer in the affirmative would resolve Tuza's conjecture, and in fact its above generalization that $\tau(H) \leq 2 \cdot G^{(2)}$ for all 3-uniform hypergraphs \mathcal{H} .

Problem 175. Is it true that for every 3-uniform hypergraph \mathcal{H} without a 3-tent, $\tau(\mathcal{H}) \leq 2 \cdot G(\mathcal{H})$?

We answer this question in the negative. In fact, we prove a stronger statement that there can be no forbidden substructure-based Tuza's theorem.

Theorem 176. Let $\mathcal{F} = \{F_1, F_2, \dots, F_t\}$ be an arbitrary set of t -uniform hypergraphs such that for every t -uniform hypergraph \mathcal{G} , the blown-up hypergraph $\mathcal{G}^{(t-1)}$ does not contain any $F_i \in \mathcal{F}$ as a sub-hypergraph.

Then, for every $\epsilon > 0$, there exists a hypergraph \mathcal{H}^ϵ that does not contain any member F of \mathcal{F} as a sub-hypergraph and which satisfies $\tau(\mathcal{H}^\epsilon) \leq (1 + \epsilon) \cdot G(\mathcal{H}^\epsilon)$.

By setting \mathcal{F} to be the single 3-tent hypergraph, we obtain a counterexample to Problem 175. Furthermore, when $t = 3$, the construction we give to prove Theorem 176 will belong to the class of 3-uniform hypergraphs \mathcal{H} obtained from a given graph \mathcal{G} with the vertex set of \mathcal{H} being the edge set of \mathcal{G} , and every triangle in \mathcal{G} forming an edge in \mathcal{H} . Thus, there is no "local" proof of Tuza's conjecture that uses only substructure properties of the underlying hypergraph.

We call a hypergraph non-trivial if it has at least two edges. Before we prove the above theorem, we use a definition from [FM08].

Definition 177. Let \mathcal{F} be a non-trivial t -uniform hypergraph. Then,

$$f(\mathcal{F}) = \max_{\mathcal{F}^0 \subseteq \mathcal{F}} \frac{e^0 - 1}{v^0 - t}$$

where \mathcal{F}^0 is a non-trivial subhypergraph of \mathcal{F} with $e^0 > 1$ edges and v^0 vertices.

We now return to the proof of Theorem 176.

Proof. (of Theorem 176) We will first prove that $f(F_i) > \frac{1}{t-1}$ for all $i \in [t]$. Suppose for contradiction that there exists a uniform hypergraph $F_i \in \mathcal{F}$ such that $f(F_i) \leq \frac{1}{t-1}$. Without loss of generality, we can assume that F_i is connected. Order the edges of F_i as e_1, e_2, \dots, e_m such that for every $j > 1$, $e_j \setminus (e_1 \cup e_2 \cup \dots \cup e_{j-1}) \neq \emptyset$. For every $j \in [m]$; let F_j^0 be the subhypergraph induced by e_1, e_2, \dots, e_j . As $f(F_i) \leq \frac{1}{t-1}$, we can infer that for every $j > 1$,

$$\frac{|E(F_j^0)| - 1}{|V(F_j^0)| - t} \leq \frac{1}{t-1}$$

which implies that $|V(F_j^0)| \leq (t-1)|E(F_j^0)| + 1 = (t-1)j + 1$. As $|V(F_j^0)| = |V(F_{j-1}^0)| + |e_j \setminus (e_1 \cup e_2 \cup \dots \cup e_{j-1})|$, we get that $|V(F_j^0)| \leq |V(F_{j-1}^0)| + t - 1$, which combined with the above shows that the inequality is in fact tight for every $j \in [m]$. Thus, for every $j > 1$; $|V(F_j^0)| = |V(F_{j-1}^0)| + t - 1$, which implies that for every $j > 1$,

$$|e_j \setminus (e_1 \cup e_2 \cup \dots \cup e_{j-1})| = 1:$$

We now construct a t -uniform hypergraph \mathcal{H} such that F_i is isomorphic to $\mathcal{H}^{(t-1)}$. We construct the hypergraph \mathcal{H} inductively via $\mathcal{H}_1, \mathcal{H}_2, \dots, \mathcal{H}_m = \mathcal{H}$ such that $\mathcal{H}_j^{(t-1)}$ is isomorphic

to F_j^0 for all $j \geq 2$ [m]. First, we set the hypergraph H_1 to be equal to the uniform hypergraph on t vertices with a single edge e_1 is trivially isomorphic to F_1^0 . Assume by inductive hypothesis that there is a hypergraph H_k such that F_k^0 is isomorphic to $H_k^{(t-1)}$ for some $k \geq 2$ [m-1]. Let $f: F_k^0 \rightarrow H_k^{(t-1)}$ be the isomorphism between the two hypergraphs. The hypergraph F_{k+1}^0 obtained from F_k^0 by adding an edge e_{k+1} such that e_{k+1} intersects with F_k^0 in exactly one vertex $v \in V(F_k^0)$. Recall that the vertex set of $H_k^{(t-1)}$ is the set of subsets of vertices of H_k of size $t-1$. Thus, $v = \{p_1, p_2, \dots, p_{t-1}\}$ for a set of vertices $p_1, p_2, \dots, p_{t-1} \in V(H_k)$. We construct H_{k+1} by introducing a new vertex x and adding the edge $e_{k+1} = \{x, p_1, p_2, \dots, p_{t-1}\}$ to the hypergraph H_k . Thus, $H_{k+1}^{(t-1)}$ is obtained from $H_k^{(t-1)}$ by adding single edge that intersects with $H_k^{(t-1)}$ at exactly one vertex, that is $\{p_1, p_2, \dots, p_{t-1}\}$. Hence, $H_{k+1}^{(t-1)}$ is isomorphic to F_{k+1}^0 , completing the proof. This proves that there exists a uniform hypergraph $H = H_m$ such that $F_i = H^{(t-1)}$, contradicting the fact that $n^{(t-1)}$ -blown-up hypergraph contains F as a subhypergraph.

Thus, $(F_i) > \frac{1}{t-1}$ for all $i \geq 2$ [t]. Let $\epsilon = \min_{i \geq 2} (F_i) > \frac{1}{t-1}$. Consider a random uniform hypergraph H^0 on n vertices sampled by picking each edge independently with probability $p = n^{-\frac{1}{t-1}}$. We now delete the edges in a maximal collection of edge disjoint copies of members of F from H^0 . It has been proved [BFM10; FM08] that the maximum independent set of this construction satisfies

$$|I(H^0)| \leq n^{\frac{1}{t-1}}$$

with high probability. Thus, there exists a uniform hypergraph H^0 with n vertices without any substructure from F such that $|I(H^0)| \leq (1 - \epsilon)n$. Since for any uniform hypergraph H^0 on n vertices, $|I(H^0)| \geq \frac{1}{t}n$, this proves the claimed factor $(t-1)$ gap between $|I(H^0)|$ and $|I(H^0)|$ for every positive constant $\epsilon > 0$. □

9.4.1 Explicit construction of tent-free hypergraphs

We now describe an explicit hypergraph giving negative answer to Problem 175. Our counterexample is a hypergraph with vertex set $[3]^n$ for large enough n and the edge set is the set of all combinatorial lines that we formally define below:

Definition 178. (Combinatorial lines in $[3]^n$) A set of three distinct vectors $u = (u_1; u_2; \dots; u_n); v = (v_1; v_2; \dots; v_n); w = (w_1; w_2; \dots; w_n) \in [3]^n$ forms a combinatorial line if there exists a subset $S \subseteq [n]$ such that

1. For all $i \in [n] \setminus S$, $u_i = v_i = w_i$.
2. There exist three distinct integers $a, v^0, w^0 \in [3]$ such that for all $i \in S$, $u_i = a, v_i = v^0, w_i = w^0$.

We will use the following seminal result about combinatorial lines:

Theorem 179. (Density Hales Jewett Theorem [FK91],[Pol12]) For every positive integer k and every real number $\epsilon > 0$ there exists a positive integer $D = D(k, \epsilon)$ such that if $n \geq D$ and A is any subset of $[k]^n$ of density at least ϵ , then A contains a combinatorial line.

We now give a proof of Theorem 163.

Proof. The hypergraph that we use $\mathcal{H} = (V; E)$ has $V = [3]^n$ for n large enough to be set later, and the edges are all the combinatorial lines of $[3]^n$. First, we claim that the above defined hypergraph does not have a $\frac{2}{3}$ -tent. Suppose for contradiction that there are edges e_1, e_2, e_3, e_4 satisfying the properties of Definition 174. Let $u = (u_1; u_2; \dots; u_n) \in e_4 \setminus e_1; v = (v_1; v_2; \dots; v_n) \in e_4 \setminus e_2; w = (w_1; w_2; \dots; w_n) \in e_4 \setminus e_3$. Note that $e_4 = \{u; v; w\}$. Thus, there exists a subset $S \subseteq [n]$ such that for all $i \in [n] \setminus S, u_i = v_i = w_i$. Without loss of generality, we can also assume that for all $i \in S, u_i = 1; v_i = 2; w_i = 3$.

Let $x = (x_1; x_2; \dots; x_n) \in e_1 \setminus e_2 \setminus e_3$. Note that $\{x; u\} \subseteq e_1; \{x; v\} \subseteq e_2; \{x; w\} \subseteq e_3$. Consider an arbitrary element $p \in S$, and without loss of generality, let $x_p = 1$. Thus, we have that $x_p = 1; v_p = 2$ and both $x; v$ share the combinatorial line e . This implies that there exist a subset $S_2 \subseteq [n]$ such that for all $i \in [n] \setminus S_2, x_i = v_i$ and for all $i \in S_2, x_i = 1; v_i = 2$. Similarly, there exists a subset $S_3 \subseteq [n]$ such that for all $i \in [n] \setminus S_3, x_i = w_i$ and for all $i \in S_3, x_i = 1; w_i = 3$.

Note that $S_2 \subseteq S$. Suppose for contradiction that there exists $S_2 \not\subseteq S$. Then, we have $v_j = 2; x_j = 1$. However, since $v_i = w_i$ for all $i \in [n] \setminus S$, we get that $w_j = 2$, and thus $j \notin S_3$, which implies that $x_j = w_j = 2$, a contradiction. Thus $S_2 \subseteq S$, and similarly $S_3 \subseteq S$. We can also observe that $S \subseteq S$ since in that case $x = u$ which cannot happen since $e_4 \setminus e_2 = \{u\}$. By the same argument on e_3 , we can deduce that $S \subseteq S$. As S_2 is a strict subset of S , there exists $j \in S \setminus S_2$. As $v_i = x_i$ for all $i \in [n] \setminus S_2, x_j = v_j = 2$. As $j \in S$, we have $w_j = 3$. However, as $w_j \notin x_j$, this implies that $j \notin S_3$, which then implies that $x_j = 1$, a contradiction.

Now, we will prove that for large enough $n, (H) > (3 - \frac{1}{3}) (H)$. Let $N = 3^n$. Since the cardinality of V is equal to N , we have $(H) = \frac{N}{3}$. We apply Theorem 179 with $k = 3; \epsilon = \frac{1}{3}$, and set $n = \text{DHJ}(k; \epsilon)$. Thus, we can infer that in any subset $T \subseteq V$ of size $\frac{2}{3}N$, there exists an edge e fully contained in T . Thus, we get that $(H) > (1 - \frac{1}{3})N$, which gives $(H) > (3 - \frac{1}{3}) (H)$. \square

9.5 Vertex cover and set cover on simple hypergraphs

As mentioned earlier, the edges in a $(t-1)$ -blown-up hypergraph of a uniform hypergraph can intersect on at most one element, so such hypergraphs are simple. In this section, we will take a step back and address to what extent improved approximation algorithms are possible for vertex cover on simple hypergraphs. We will also consider the dual problem, of covering the vertices by the fewest possible hyperedges, namely the set cover problem, on simple hypergraphs but without any restriction on the size of the hyperedges. Note that a hypergraph is simple if and only if the edge-vertex incidence bipartite graph does not contain a copy of $K_{2,t}$. Thus, a hypergraph is simple if and only if its dual is simple.

9.5.1 Vertex cover on simple t -uniform hypergraphs

We now prove Theorem 164 which shows that simple hypergraphs are still rich enough to preclude a non-trivial approximation to vertex cover. Our hardness is established using a reduction from

the general problem of vertex cover on uniform hypergraphs. In particular, we use the following result:

Theorem 180. ([Din+05]) For every constant $\epsilon > 0$ and $t \geq 3$, the following holds: Given a t -uniform hypergraph $G = (V; E)$, it is NP-hard to distinguish between the following cases:

1. **Completeness:** G has a vertex cover of measure $1 - \epsilon$.
2. **Soundness:** Any subset S of measure ϵ contains an edge from E .

We give a randomized reduction from Theorem 180 to Theorem 164. The approach is similar to the one used in [GL17] for showing the inapproximability of Transversal in graphs. It was also used in the recent tight hardness for Max Coverage on simple set systems [CKL21].

Let us instantiate Theorem 180 with ϵ replaced by $\epsilon^0 = \frac{\epsilon}{4}$, and let the resulting hypergraph be denoted by G . Now, given this t -uniform hypergraph $G = (V; E)$, we output a t -uniform hypergraph $H = (V^0; E^0)$ as follows: Let $n = |V|$; $m = |E|$. We have integer parameters $B; P$ depending on $t; n; m$ to be set later. The vertex set V^0 is $V \cup [B]$ —we have a cloud of B vertices $v^1; v^2; \dots; v^B$ in V^0 corresponding to every vertex $v \in V$. For every edge $e = (v_1; v_2; \dots; v_t) \in E$, we pick P edges $e^1; e^2; \dots; e^P$ with $e^j = ((v_1)_i; (v_2)_i; \dots; (v_t)_i)$ and add them to E^0 , where for each $i \in [t]$ and $j \in [P]$, $(v_j)_i$ is chosen uniformly and independently at random from $\{v_j^1; v_j^2; \dots; v_j^B\}$. Thus, so far, we have added mP edges to E^0 .

We first upper bound the expected value of the number of pairs of edges that intersect in more than one vertex. Order the edges in E^0 as $e_1; e_2; \dots; e_{mP}$. Let X denote the random variable that counts the number of pairs of edges in E^0 that intersect in more than one vertex. For every pair of indices $i; j \in [mP]$, let the random variable X_{ij} be the indicator variable of the event that the edges e_i and e_j of E^0 intersect in greater than one vertex. Note that the edges e_i corresponding to e_j and e_j have at most t vertices in common. Thus, the probability that e_i and e_j intersect in at least two vertices is upper bounded by $\frac{1}{B^2}$. Summing over all the pairs $i < j$, we get

$$E[X] \leq \binom{mP}{2} \frac{1}{B^2} = \frac{m^2 t^2 P^2}{2B^2}.$$

By Markov's inequality, with probability at least $\frac{9}{10}$, X is at most $\frac{10m^2 t^2 P^2}{B^2}$.

We consider all the pairs of edges that intersect in more than one vertex and arbitrarily delete one of those edges. Let the resulting set of edges be denoted by E^{00} . The hypergraph resulting in this reduction is $H = (V^0; E^{00})$. Note that H is indeed a simple hypergraph. We will prove the following:

1. **(Completeness)** If G has a vertex cover of measure $1 - \epsilon$, then there is a vertex cover of measure $1 - \epsilon^0$ in H .
2. **(Soundness)** If every subset S of measure ϵ^0 contains an edge from E , then with probability at least $\frac{9}{10}$, every subset S^0 of V^0 of measure ϵ contains an edge from E^{00} .

Completeness. If G has a vertex cover of size $(1 - \epsilon)n$, then picking all the vertices in V^0 in the cloud corresponding to these vertices ensures that H has a vertex cover of size $(1 - \epsilon^0)n$. Thus, in the completeness case, there is a vertex cover of measure $1 - \epsilon^0$.

Soundness. Suppose that every measure subset V contains an edge from E . Our goal is to show that with probability at least $\frac{1}{5}$, every measure subset V^0 contains an edge from E^0 . We first prove the following lemma:

Lemma 181. With probability at least $\frac{9}{10}$ over the choice of E^0 , the following holds: For every edge $e = (v_1; v_2; \dots; v_t) \in E$, and every subset $S \subseteq V^0$ such that for each $i \in [t]$, S contains at least $\frac{1}{4}B$ vertices from $v_i^1; v_i^2; \dots; v_i^B$, there exists an edge $e' \in E^0$ all of whose vertices are in S .

Proof. The probability that there exists an edge $e = (v_1; v_2; \dots; v_t)$ and a subset S which contains at least $\frac{1}{4}B$ vertices from each cloud and does not contain any edge from E^0 is at most

$$m 2^{tB} \left(1 - \frac{1}{4}\right)^t \leq m 2^{tB} \log e \frac{t^P}{4^t} \leq \frac{1}{10}$$

when $P = m a B$ where $a := a(t) = \frac{4^{t+2} t}{t}$. □

Using the above lemma, we can conclude that with probability at least $\frac{10m^2 t^2 P^2}{B^2} = 10m^4 t^2 a^2$ and for every edge $e \in E$ and every subset $S \subseteq V^0$ such that for each $i \in [t]$, S contains at least $\frac{1}{4}B$ vertices from $v_i^1; v_i^2; \dots; v_i^B$, there exists an edge $e' \in E^0$ all of whose vertices are in S . We claim that this implies that with probability at least $\frac{1}{5}$, every measure subset V^0 contains an edge from E^0 . Consider an arbitrary subset V^0 such that $|U_j| \leq \frac{nB}{2}$. We choose B large enough such that $\frac{10m^4 t^2 a^2}{2} \geq \frac{nB}{2}$. Thus, the set of the vertices W in the edges deleted from E^0 to obtain E^0 has cardinality at most $\frac{nB}{2}$.

Let $U^0 = U \cap W$. Note that all the edges in U^0 that are in E^0 are present in E^0 as well. As U^0 has a measure of at least $\frac{1}{2}$ in V^0 , for at least $\frac{1}{4}n$ vertices v in V , U^0 should contain at least a fraction of the vertices in the cloud $v^1; v^2; \dots; v^B$. Since otherwise, the cardinality of U^0 is at most $n \left(\frac{n}{4} \frac{B}{4} + \frac{n}{4} B\right) < \frac{nB}{2}$, a contradiction. By Lemma 181, we can deduce that there exists an edge $e \in E^0$ all of whose vertices are in U^0 , which implies that the edge e is in E^0 as well. This proves that in the soundness case, with probability at least $\frac{1}{5}$, there exists an edge in every measure subset V^0 .

This completes the proof of Theorem 164. Under the Unique Games Conjecture [Kho02a], the hardness of vertex cover in uniform hypergraphs can be improved to $\frac{1}{2}$. We remark that we can get the same hardness for simple hypergraphs by our reduction.

9.5.2 Set Cover on Simple Set Systems

In the set cover problem, there is a set family $\mathcal{S} \subseteq 2^X$ on a universe $X = [n]$, and the goal is to cover the universe $[n]$ with as few sets from the family as possible. The greedy algorithm where we repeatedly pick the set that covers the maximum number of new elements achieves a $\ln n$ -factor approximation algorithm for the problem, and this is known to be optimal. We consider the same problem under the restriction that the family \mathcal{S} is a simple set system i.e. for every $i \in [n]$, $|S_i \cap S_j| \leq 1$. Surprisingly, in contrast with the hardness result for vertex cover, simplicity of the set family helps in achieving better approximation factor for the set cover problem.

Theorem 182. (Theorem 165 restated) The greedy algorithm achieves a $1 + \frac{\ln n}{2}$ -approximation guarantee for the set cover problem on simple set systems over a universe of size n . Furthermore, the bound is essentially tight for the greedy algorithm—there is a simple set system on which the approximation factor of greedy exceeds $(\frac{\ln n}{2} - 1)$.

Proof. First, we prove the upper bound. Let the optimal solution size be equal to k . There is $T = \{S_1, S_2, \dots, S_k\} \subseteq \mathcal{S}$ such that the union of sets in T is equal to $[n]$. For every set $S \in \mathcal{S} \setminus T$, $|S \cap T^c| \leq k$ by the simplicity of the set system, and thus, we get that

$$|S| \leq k + |S \cap T| \leq k + k = 2k \quad (9.1)$$

We now consider two different cases:

1. Suppose that $k \geq \frac{n}{2}$. We recall that the greedy algorithm in fact achieves a $1 + \frac{\ln n}{2}$ -factor approximation algorithm for set cover on general instances where n_j is the size of the largest set in the family. Thus, after the greedy algorithm picks each of which cover at least $\frac{n}{2}$ new elements, in the remaining instance, we have $n_j \leq \frac{n}{2}$. As there are k sets that cover the remaining instance, the greedy algorithm picks at most k sets after picking the sets. As each of the sets cover at least $\frac{n}{2}$ new elements, $\frac{n}{2}$. Overall, the total number of sets used by the greedy algorithm is equal to

$$t + k \frac{\ln n}{2} \leq \frac{n}{2} + k \frac{\ln n}{2} = (1 + \frac{\ln n}{2})k$$

2. Suppose that $k < \frac{n}{2}$. In this case, using (9.1), we can infer that there are at most k sets with size at least k in the family. Thus, after the greedy algorithm picks k sets, in the remaining instance, each set has size at most k and thus, greedy algorithm picks at most $k \ln k$ sets. Overall, the total number of sets picked by the greedy algorithm is equal to

$$k + k \ln k \leq k + k \frac{\ln n}{2} = (1 + \frac{\ln n}{2})k$$

Thus, in both the cases, the greedy algorithm picks at most $(1 + \frac{\ln n}{2})k$ sets.

A hard instance for the greedy algorithm. We now prove that the above bound is tight for the greedy algorithm. Fix a large integer k , and let $n = k^2$, $X = [n]$. We first add k sets to the family $\mathcal{S} = \{S_1, S_2, \dots, S_k\}$ where $S_j = \{(j-1)k+1, (j-1)k+2, \dots, jk\}$. Note that these k sets together cover the whole universe X . We view the universe X as k blocks, with the j 'th block comprising of the set S_j .

We now add $m = (k-1) \ln k$ additional pairwise disjoint sets T_1, T_2, \dots, T_m to \mathcal{S} such that the greedy algorithm picks the set T_i in the i 'th iteration. We choose the sets $T_i, i \in [m]$, as follows:

1. For $j \in [k]$, let the set X_j be the uncovered elements of the block S_j . Let $a_j = |X_j|$ for all $j \in [k]$. We initially set $X_j = S_j$ for all $j \in [k]$.
2. At every iteration $i \in [m]$:

(a) Sort the elements a_1, a_2, \dots, a_k such that $a_1 \leq a_2 \leq \dots \leq a_k$ where $(1, 2, \dots, k)$ is a permutation of $[k]$. Let $p = a_1 \leq k$ be the largest number of uncovered elements in a block.

(b) Let $P = \{1, 2, \dots, p\}$. For $l \in [p]$, let $u_l \in [n]$ be equal to the largest element in X_{p-l} .

$$u_l = \max\{b_j \mid b_j \in X_{p-l}\}$$

We set

$$T_i = \{u_l \mid l \in [p]\}$$

Furthermore, we set $X_{p-l} = X_{p-l} \setminus u_l$ for all $l \in [p]$. We also update $a_j = j$ for all $j \in [k]$.

In the above procedure to output the sets T_i , the cardinality of T_i is $|T_i| = |T_{i+1}| + p$ for all i . Furthermore, in the i th iteration of the above procedure, the cardinality of T_i is at least the number of elements in any block that are not covered yet. This ensures that the greedy algorithm in the i th iteration picks the set T_i . Furthermore, as the sets T_i are all mutually disjoint, and intersect each block at most once, the resulting set system is indeed a simple set system.

Our goal is to prove that after all the sets are picked, there are still uncovered elements in $[n]$. For an integer $i \in [m]$, we let $s_i \in [n]$ denote the number of elements not covered by the greedy algorithm before the set T_i is picked. For $i \in [m]$, the size of the set T_i picked by the greedy algorithm in the i th iteration is equal to the largest number of uncovered elements in a block i.e. the value a_1 in the i th iteration. Based on the updating procedure followed above, we can infer that this value is equal to $\frac{s_i}{k}$. This follows from the fact that at any iteration of the above procedure, the sorted values a_1, \dots, a_k satisfy $a_1 \leq a_k + 1$.

We have $s_1 = n = k^2$, and

$$s_{i+1} = s_i - |T_i| = s_i - \frac{s_i}{k} = s_i \left(1 - \frac{1}{k}\right)$$

By setting $t_i = s_i + k$ for $i \in [m]$, we get

$$t_{i+1} = t_i \left(1 - \frac{1}{k}\right)$$

Thus, we get

$$\begin{aligned} t_{m+1} &= t_1 \left(1 - \frac{1}{k}\right)^m \\ &= (k^2 + k) \left(1 - \frac{1}{k}\right)^{(k-1) \ln k} \\ &= (k^2 + k) \exp\left(-\frac{1}{k} (k-1) \ln k\right) \quad \left(\text{Using } 1 - x \leq e^{-\frac{x}{80}} \text{ for } x < 1\right) \\ &= k + 1 \end{aligned}$$

Thus, $s_{m+1} > 1$, which proves that there are elements that are not covered after the greedy algorithm uses $m = (k-1) \ln k$ sets. This completes the proof that there are simple set systems on n elements with $k = \frac{n}{\ln n}$ sets covering all the elements where as the greedy algorithm picks at least $(k-1) \ln k = k \frac{\ln n}{2} - 1$ sets.

□

Chapter 10

Scheduling with non-uniform communication delay

10.1 Introduction

We study the problem of scheduling jobs with precedence and non-uniform communication delay constraints on identical machines to minimize the makespan objective function. This classic model was first introduced by Rayward-Smith [Ray87] and Papadimitriou and Yannakakis [PY90]. In this problem, we are given a set of n jobs, where each job j has a processing length $p_j \in \mathbb{Z}_+$. The jobs need to be scheduled on m identical machines. The jobs have precedence and communication delay constraints, which are given by a partial order constraint $j \prec j^0$ encodes that job j^0 can only start after job j is completed. Moreover, if $j \prec j^0$ and j, j^0 are scheduled on different machines, then j^0 can only start executing at least c_{jk} time units after j had finished. On the other hand, if j and j^0 are scheduled on the same machine, then j^0 can start executing immediately after j finishes. The goal is to schedule jobs non-preemptively to minimize the makespan objective function, which is defined as the completion time of the last job. In a non-preemptive schedule, each job needs to be assigned to a single machine and executed during a contiguous time interval of length p_j . In the classical scheduling notation, the problem is denoted by $P | j \text{ prec}, c_{jk} | C_{\max}$.¹ A closely related problem is $B1 | j \text{ prec}, c_{jk} | C_{\max}$, where the scheduler has access to an unbounded number of machines.

Scheduling jobs with precedence and communication delays has been studied extensively over many years [Ray87; PY90; MK97; HM01; TY92; HLV94; Gir+08]. Furthermore, due to its relevance in datacenter scheduling problems and large-scale training of ML models, there has been a renewed interest in more applied communities; we refer the readers to [Cho+11; Guo+12; HCG12; Shy+18; Zha+12; Zha+15; Luo+16; Nar+19; Mir+17; GCL18; JZA19; Tar+20]. However, from a theoretical standpoint, besides NP-hardness results, very little was known in

¹We adopt the convention of [Gra+79; VLL90], where the respective fields denote (1) machine environment: Q for related machines, P for identical machines; (2) job properties: prec for precedence constraints, c_{jk} for communication delays, when all the communication delays are equal, $c_{jk} = c$; $p_j = 1$ for the unit-length case; (3) objective: C_{\max} for minimizing makespan.

terms of the algorithms for the problem until the recent work by Maiti et al. [Mai+20] and Davies et al. [Dav+20; Dav+21]. These very recent papers designed polylogarithmic approximation algorithms for the special case when all the communication delays are equal. We survey these results in Section 10.1.2. In fact, the problems of scheduling jobs with communication delays are some of the well-known open questions in approximation algorithms and scheduling theory, and have resisted progress for a long time. For this reason, the influential survey by Schuurman and Woeginger [SW99b] and its recent update by Bansal [Ban17] list understanding the approximability of the problems in this space as one of the top-10 open questions in scheduling theory.

In particular, an open problem in these surveys asks if the non-uniform communication delay problem on identical machines, even assuming an unbounded number of machines ($\forall p \leq c_k \leq C_{\max}$), admits a constant-factor approximation algorithm. We answer this question in the negative.

Theorem 183. For every constant $\epsilon > 0$, assuming $\text{NP} \neq \text{ZTIME}_{n^{(\log n)^{O(1)}}}$, the non-uniform communication delay problem ($\forall p \leq c_k \leq C_{\max}$) does not admit a polynomial-time $(\log n)^{1/\epsilon}$ -approximation algorithm.

We remark that our hard instances contain only two distinct values of communication delays (essentially 0 and 1). Furthermore, for $\forall p \leq c_k \leq C_{\max}$, the problem with an unbounded number of machines, is a special case of $\forall p \leq c_k \leq C_{\max}$, where the number of machines is specified, our theorem also implies the same hardness for $\forall p \leq c_k \leq C_{\max}$.

10.1.1 Our Techniques

Our hardness result is obtained using a reduction via a problem we call Unique Machines Precedence constraints Scheduling (UMPS). In this problem, there are m machines and jobs j_1, j_2, \dots, j_n with precedence relations between them. Each job j_i has length $p(i)$ and can be scheduled only on a unique machine $M(i) \in [m]$. The objective is to schedule the jobs non-preemptively on the corresponding unique machines, respecting the precedence relations, so as to minimize the makespan objective function. Our proof of Theorem 183 proceeds via two steps:

1. First we show a reduction from an instance of the UMPS problem to an instance of the non-uniform communication delay problem. The key step is to make sure that the set of jobs $J(i)$ that need to be scheduled on machine i do not get scheduled on multiple machines in \mathcal{J} . We achieve this by introducing a dummy job and introducing precedence constraints from all jobs $J(i)$ to j_i and a very large communication delay. This ensures that $J(i)$ and j_i are scheduled on the same machine in \mathcal{J} , although this machine need not be i . Despite this, we show that any valid schedule \mathcal{J} can be mapped back to a feasible schedule of \mathcal{I} , with almost the same makespan. Our reduction creates only two types of communication delays and works for the unit-length case.
2. Next we observe that the UMPS problem generalizes the classical job shops problem (see e.g. [Law+93; LMR94; MS11]), whose approximation is well understood [SSW94; CS00; Gol+01; FS02]. The logarithmic hardness result for the acyclic job shops problem by




Figure 10.1: Role of the UMPS problem in our hardness reduction.

Mastrolilli and Svensson [MS11] implies a logarithmic hardness of the UMPS problem. We remark that the hardness result of [MS11] only works when jobs have lengths, and hence our Theorem 183 only applies to the setting where jobs have lengths.

In hindsight, our proof of Theorem 183 is surprisingly simple. However, the main conceptual contribution of our proof is in identifying the UMPS problem as a central problem that has implications for the hardness of various scheduling problems. Furthermore, the UMPS problem, which can be viewed as a generalization of the job shop scheduling model or as a highly restricted version of multidimensional scheduling with precedences, or as a restricted assignment problem with precedence constraints, is a fundamental problem to study on its own, both from a theoretical perspective and also from a practical point of view. We believe the UMPS problem is a key intermediate step towards resolving several long-standing open problems in scheduling theory. We make the following two conjectures regarding the approximability of UMPS.

Conjecture 184. There exists a constant $\epsilon < 1$ such that it is NP-hard to approximate UMPS within a factor of $1/\epsilon$, even when all jobs have unit lengths, where n is the number of jobs.

Conjecture 185. There exists an absolute constant $\epsilon < 1$ such that the following holds. For every constant $\delta > 0$, it is NP-hard to approximate UMPS within a $(\log n)^{1/\delta}$ -factor, even when the number of machines m is at most $(\log n)^C$ and all the jobs have unit lengths, where n is the number of jobs.

Our second main contribution is to show that the above conjectures imply hardness results for various problems. In particular, Conjecture 185 implies logarithmic hardness for scheduling with precedences on related machines, another top-10 problem in scheduling theory [SW99b; Ban17] and in the approximation algorithm book of Shmoys and Williamson [WS11].

Theorem 186. Assuming Conjecture 185 and $\text{NP}^* \text{DTIME } n^{(\log n)^{O(1)}}$, there exists an absolute constant $\delta > 0$ such that the problem of scheduling related machines with precedences $(Q_j \text{ prec}_j C_{\max})$ has no polynomial-time $(\log m)^{\delta}$ -factor approximation algorithm.

Previously, Bazzi and Norouzi-Fard [BN15] introduced a partite hypergraph partition problem whose hardness implies a superconstant hardness for scheduling with precedences on related machines. Our reduction uses the same idea of job replication as [BN15], while our soundness analysis is technically more involved. We also show that the hypothesis of [BN15] implies a superconstant hardness of the UMPS problem. Thus, our problem can be viewed as a weaker version of the hypothesis of [BN15] with the same implication towards the hardness of related machines. Furthermore, stronger hardness of the UMPS problem implies better (almost optimal) hardness results for the related machines scheduling problem.

Finally, we note that Conjecture 184 implies that precedence-constrained scheduling (even

without communication delays) is very hard to approximate when generalized to the restricted assignment setting or unrelated machines.

Our confidence in the above conjectures stems from the fact that existing techniques, both the classical jobshops algorithms [LMR94] and the recent LP-hierarchies-based algorithms [Mai+20; Dav+20] fail to give non-trivial approximation guarantees for the UMPS problem. Furthermore, a candidate hard instance for the problem is a layered instance, where there are precedences between jobs $j_1 \prec j_2$ only if j_1 can be scheduled on the machine i and j_2 can be scheduled on the machine $i + 1$. These layered instances are closely related to the partite partitioning hypothesis of [BN15] and the integrality gap instances [Mai+20] for the problem of scheduling with uniform communication delays.

10.1.2 A Brief History of the Communication Delay Problem

In this subsection, we give a brief overview of the literature on the problem of scheduling with communication delays.

Scheduling with precedences. Scheduling with precedences to minimize makespan $P|prec|C_{max}$ is a classic combinatorial optimization problem and is a special case of the communication delay problem with $c = 0$ for all pairs of jobs. In one of the earliest results in the scheduling theory, Graham's list scheduling algorithm [Gra66] was shown to be a 2 -factor approximation for the problem. Recently, Svensson [Sve10] gave a matching hardness of approximation result assuming (a variant of) the Unique Games Conjecture [BK09]. When the number of machines is a constant, a series of recent works have obtained $(c+1)$ -approximation in nearly quasi-polynomial time [LR16; Gar18; Kul+20; Li21].

Uniform communication delay setting. The problem becomes much harder with communication delays, even when all the communication delays are equal. This problem is denoted by $P|j\text{ prec }c|C_{max}$ and is referred to as scheduling with uniform communication delays. In this setting, Graham's list scheduling algorithm obtains a $(c+1)$ -factor approximation. This was improved to $2=3(c+1)$ by Giroudeau et al. [Gir+08] in the case when the jobs have unit lengths $P|j\text{ prec }p_j=1;c=2|C_{max}$. In recent concurrent and independent works, poly-logarithmic-factor approximation algorithms have been obtained for the uniform communication delays problem $P|j\text{ prec }c|C_{max}$ by Maiti et al. [Mai+20] and Davies et al. [Dav+20; Dav+21].

On the hardness front, when $c=1$, Hoogeveen, Lenstra and Veltman [HLV94] showed that the problem $P|j\text{ prec }p_j=1;c=1|C_{max}$ is NP-hard to approximate to a factor better than $7/6$. The result has been generalized for 2 to $(1+1/(c+4))$ -hardness [Gir+08].

Scheduling with non-uniform communication delay. We do not know of any algorithm for the non-uniform communication delay $P|j\text{ prec }c_{jk}|C_{max}$ problem. On the hardness side,

²Papadimitriou and Yannakakis [PY90] claim the hardness for $P|j\text{ prec }p_j=1;c|C_{max}$, but give no proof. Schuurman and Woeginger [SW99b] remark that "it would be nice to have a proof for this claim".

the best hardness known is the above small constant hardness of the uniform communication delay setting. While our main result shows logarithmic hardness for this problem, it is conceivable that it admits a polylog-approximation algorithm, although our conjectures suggest otherwise.

Duplication model. Scheduling with communication delays problem has also been studied in the duplication model, where we allow jobs to be duplicated (replicated), i.e., executed on more than one machine to avoid communication delays. In this easier model, for the general $P1 \mid j \text{ prec } p_j; c_k; \text{dup } j \mid C_{\max}$ problem, there is a simple factor approximation algorithm by Papadimitriou and Yannakakis [PY90]. On the other hand, [PY90] also show the NP-hardness of $P1 \mid j \text{ prec } p_j = 1; c; \text{dup } j \mid C_{\max}$. Note that the $O(1)$ -approximation algorithm for the version with duplication is in sharp contrast to our hardness result (Theorem 183) illustrating that the problem is significantly harder without duplication.

10.1.3 Discussion and Open Problems

While we make progress on the hardness of approximation of scheduling with non-uniform communication delay, the main conceptual contribution of this work is initiating the formal study of the UMPS problem. When jobs have lengths, the problem does not admit a polylogarithmic approximation. However, much less is known for the unit-length case. We now mention a few open problems in this direction.

1. The key open problem is to prove (or disprove) Conjecture 185. A positive resolution of the conjecture would prove the hardness of scheduling related machines with precedences, a long-standing open problem in scheduling theory. By the same reduction as in the proof of Theorem 183, Conjecture 185 also implies a logarithmic hardness of approximation for the non-uniform communication delay problem even when the jobs have unit lengths ($P1 \mid j \text{ prec } p_j = 1; c_k \mid C_{\max}$).
2. On the other hand, obtaining good approximation algorithms for the UMPS problem would be even more exciting. Is Conjecture 184 true, or is there a polylog-factor approximation algorithm for the unit-length case?

10.1.4 Organization

The rest of the chapter is organized as follows. We first formally define the UMPS problem and relate it to the jobshops problem in Section 10.2. We then use the hardness of the UMPS problem to prove Theorem 183 in Section 10.3. Finally, in Section 10.4, we show that Conjecture 185 implies an improved hardness of related machine scheduling with precedences and that the hypothesis of [BN15] implies a superconstant hardness of the UMPS problem with unit lengths.

10.2 Unique Machine Precedence Constraints Scheduling problem

We first formally define the Unique Machine Precedence constraint Scheduling (UMPS) problem.

Definition 187. (Unique Machine Precedence constraint Scheduling) In the Unique Machine Precedence constraint Scheduling (UMPS) problem, the input is a set of machines and jobs j_1, j_2, \dots, j_n with precedence relations between them. Furthermore, each job can be scheduled only on a fixed machine $M(i) \in [m]$, and takes $p(i)$ time to complete. The jobs should be scheduled non-preemptively, i.e., once a machine starts processing a job, it must finish it before processing other jobs. The objective is to schedule the jobs on the corresponding machines in this non-preemptive manner while respecting the precedence relations, so as to minimize the makespan.

We note that the UMPS problem is a generalization of the classical jobshops problem that we formally define below.

Definition 188. (Job shops) In the jobshops problem, the input is a set of jobs to be processed on a set M of m machines. Each job consists of j operations $O_{1;j}, O_{2;j}, \dots, O_{j;j}$. Operation $O_{i;j}$ must be processed for $p_{i;j}$ units of time without interruptions on the machine $m_{i;j} \in M$, and can only be scheduled if all the preceding operations $O_{s;j}, s < i$ have finished processing. The objective is to schedule all the operations on the corresponding machines to minimize the makespan.

Note that jobshops problem is a special case of the UMPS problem, corresponding to the case when the precedence DAG is a disjoint union of chains. The jobshops problem has received a lot of attention and played an important role in the development of key algorithmic techniques [Law+93; LMR94]. On the hardness front, Mastrolilli and Svensson showed almost optimal hardness results for the problem in a breakthrough result [MS11].

Theorem 189. For every constant $\epsilon > 0$, assuming $\text{NP}^* \subseteq \text{ZTIME}(n^{(\log n)^{O(1)}})$, there is no polynomial-time $(\log n)^{1/\epsilon}$ -factor approximation algorithm for the jobshops problem, where n is the total number of operations in the given jobshops instance.

As a corollary, we obtain the following hardness result.

Corollary 190. For every constant $\epsilon > 0$, assuming $\text{NP}^* \subseteq \text{ZTIME}(n^{(\log n)^{O(1)}})$, there is no polynomial-time $(\log n)^{1/\epsilon}$ -factor approximation algorithm for the UMPS problem.

10.3 Hardness of Scheduling With Non-Uniform Communication Delays

We now give a reduction from the UMPS problem to the non-uniform communication delay problem, thereby proving the hardness of the non-uniform communication delay problem. We restate the theorem for convenience.

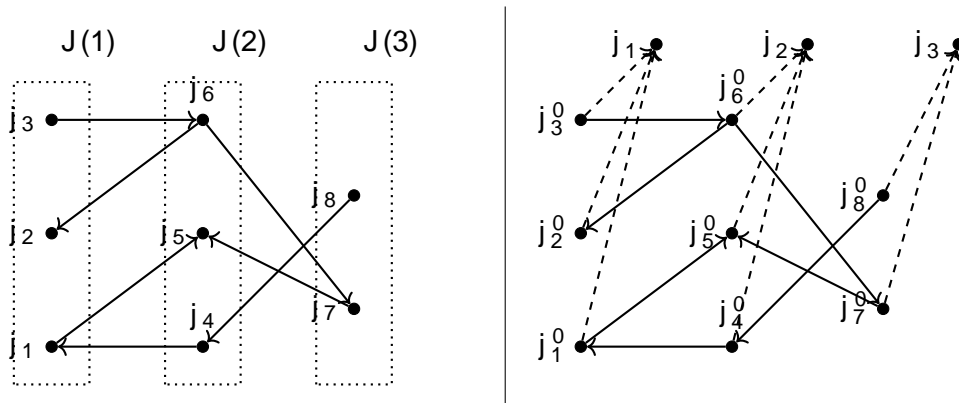


Figure 10.2: Illustration of the reduction from UMPS to non-uniform communication delays. In the communication delay instance on the right, the dashed arrow precedences have communication delay C_1 while the normal arrow precedences have communication delay 0.

Theorem 183. For every constant $\epsilon > 0$, assuming $\text{NP}^* \subseteq \text{ZTIME } n^{(\log n)^{O(1)}}$, the non-uniform communication delay problem $(P \mid j \text{ prec } c_k \mid C_{\max})$ does not admit a polynomial-time $(\log n)^{1/\epsilon}$ -approximation algorithm.

Reduction. Let I be an instance of the UMPS problem with jobs j_1, j_2, \dots, j_n , and m machines. Furthermore, each job j_i has a processing time $p(i)$ and can be scheduled only on the machine $M(i) \in [m]$. For an index $i \in [m]$, let $J(i) = \{j_1, j_2, \dots, j_n\}$ denote the set of jobs that can be scheduled on the machine i .

Roughly speaking, our idea in the reduction is to output a non-uniform communication delay instance where we force the jobs $J(i)$ to be scheduled on the same machine, for every $i \in [m]$. We achieve this by adding a set of dummy jobs $j_1^0, j_2^0, \dots, j_m^0$ and adding precedences with very large communication delay from all the jobs $J(i)$ to j_i^0 for every $i \in [m]$. More formally, we define an instance I^0 of the non-uniform communication delay problem as follows. First, we choose a large integer $C_1 = n \sum_{i=1}^m p(i)$. There are $n + m$ jobs in I^0 : a set of jobs $j_1^0, j_2^0, \dots, j_n^0$ such that for each $i \in [n]$, the processing time $p(i^0)$ is equal to $p(i)$, and a set $\{j_1^0, j_2^0, \dots, j_m^0\}$ of m dummy jobs, each with processing time 0. For every precedence relation $j_u \prec j_v$ in the original instance I , there is a precedence relation $j_u^0 \prec j_v^0$ in I^0 with communication delay 0. Finally, for every $i \in [m]$, and every job $j_l \in J(i)$, there is a precedence relation $j_l^0 \prec j_i^0$ with communication delay C_1 .

Completeness. Suppose that there is a schedule σ with makespan at most t . Then, we claim that there is a schedule σ^0 with makespan at most $t + 1$. We use machines $M(1), \dots, M(m)$ and schedule the jobs j_i^0 on the machine $M(i)$ in the same time slot used by the schedule σ . As the communication delay of the precedences among the jobs $j_1^0, j_2^0, \dots, j_n^0$ is zero, we can schedule these jobs using n machines with makespan at most t . Now, after all the jobs $j_1^0, j_2^0, \dots, j_n^0$ have been scheduled, we schedule the jobs j_i^0 on the machine $M(i)$, for every $i \in [m]$. As we are scheduling all the jobs $J(i)$ and j_i^0 on the same machine for every $i \in [m]$, we incur no communication delay when we are scheduling the dummy jobs, and we can schedule all the

dummy jobs j_1, j_2, \dots, j_m simultaneously in the time slot between L and $L + 1$.

Soundness Suppose that there is a schedule σ^0 with makespan at most L . Then, we claim that there is a schedule σ with makespan at most L as well.

Note that there is a trivial schedule σ^0 where we schedule each job one by one after topologically sorting them, that has a makespan of $\sum_{j=1}^n p(j)$. Thus, henceforth, we assume that $L \geq \sum_{j=1}^n p(j)$. For an index $i \in [m]$; let $J^0(i)$ be the subset of jobs in J^0 whose corresponding jobs in J are to be scheduled on the machine

$$J^0(i) := \{j \in J^0 : M(j) = i\}.$$

We claim that in the schedule σ^0 with makespan at most L , for every $i \in [m]$, all the jobs in $J^0(i)$ must be scheduled on the same machine. Suppose for the sake of contradiction that this is not the case. If there are jobs $j_{i_1}^0$ and $j_{i_2}^0$ such that $M(j_{i_1}^0) = i_1$ and $M(j_{i_2}^0) = i_2$ are scheduled on different machines i_1, i_2 in the schedule σ^0 , at least one of $j_{i_1}^0$ and $j_{i_2}^0$ is scheduled on a different machine than j_i . However, as there are precedence relations $j_{i_1}^0 \rightarrow j_i^0$ and $j_{i_2}^0 \rightarrow j_i^0$ with communication delay C_{i_1} , at least one of the precedence relations has to wait for the communication delay, and thus, the makespan is at least $L + C_{i_1}$, a contradiction.

Thus, for every $i \in [m]$, all the jobs in $J^0(i)$ are processed on the same machine i . This implies that at any point of time, at most one job from $J^0(i)$ is being processed, for every $i \in [m]$. Using this observation, we output a schedule σ for every job $j_i \in J(i)$, we schedule j_i in the same time slot used by the job j_i^0 in the schedule σ^0 . By the above observation, every machine $i \in [m]$ is used at most once at any time point. Furthermore, as the schedule σ respects the precedence conditions, the new schedule σ also respects the precedence conditions. Note that the makespan of this schedule σ is equal to L . This completes the proof that there exists a schedule σ with makespan at most L , if there exists a schedule σ^0 with makespan at most L .

This completes the proof of Theorem 183. We remark that the same reduction also proves a $(\log n)^{1/\epsilon}$ -factor inapproximability of the bounded-machines version $P_j \text{ prec } c_{jk} \leq C_{\max}$ of the non-uniform communication delay problem.

10.4 Conditional Hardness of Scheduling With Precedence Constraints on Related Machines

In this section, we first prove that Conjecture 185 implies improved hardness of scheduling related machines with precedences.

We begin by formally defining the scheduling related machines with precedences problem $(Q_j \text{ prec } C_{\max})$.

Definition 191. (Scheduling related machines with precedences) In the scheduling related machines with precedences problem, the input is a set of machines M and a set of jobs J with precedences among them. Furthermore, each machine $i \in M$ has speed $s_i \in \mathbb{Z}^+$, and each job $j \in J$ has processing time $p_j \in \mathbb{Z}^+$, and scheduling the job j on machine i takes $\frac{p_j}{s_i}$ units of time. The

objective is to schedule the jobs on the machines non-preemptively respecting the precedences constraints, to minimize the makespan.

An algorithm with $O(\log m)$ approximation guarantee for the problem was given independently by Chudak and Shmoys [CS99], and Chekuri and Bender [CB01]. On the hardness side, a hardness factor α follows from the identical machines setting [Sve10], assuming a variant of the Unique Games Conjecture. Furthermore, Bazzi and Norouzi-Fard [BN15] put forth a hypothesis on the hardness of k -partite graph partitioning problem, which implies a super constant hardness of the scheduling related machines with precedences problem.

We now prove that Conjecture 185 implies poly logarithmic hardness of scheduling related machines with precedences problem.

Theorem 186. Assuming Conjecture 185 and $\text{NP}^* \text{DTIME } n^{(\log n)^{O(1)}}$, there exists an absolute constant $\epsilon > 0$ such that the problem of scheduling related machines with precedences $(Q_j \text{ prec } C_{\max})$ has no polynomial-time $(\log m)^\epsilon$ -factor approximation algorithm.

Reduction. Our reduction is essentially the same reduction as in [BN15] where the authors obtained conditional hardness of the related machine scheduling with precedences problem assuming the hardness of a certain k -partite graph partitioning problem. However, our soundness analysis needs more technical work.

We start with an instance of the UMPS problem with m unit sized jobs j_1, j_2, \dots, j_n and m machines, and every job can only be scheduled on the machine $M(i) \in [m]$. Furthermore, we let $J(i) \subseteq [n]; i \in [m]$ denote the set of all the jobs that can be scheduled on the machine

We now output an instance I^0 of the related machine scheduling problem. We choose a parameter $\epsilon = 10n^3m$. For every $i \in [n]$, we have a set J_i of $2^{(m - M(i))}$ jobs in I^0 . The processing time of each of these jobs is equal to ϵ^{-1} . For every $i \in [m]$, we have M_i , a set of $2^{(m - i)}$ machines, each with speed ϵ^{-1} . Furthermore, for every precedence constraint $j_u \prec j_v$ in I^0 , we have $j_{i_1}^0 \prec j_{i_2}^0$ for every $j_{i_1}^0 \in J_u$ and $j_{i_2}^0 \in J_v$.

Completeness Suppose that there is a scheduling σ with makespan equal to L . Then, we claim that there is a scheduling σ^0 with makespan at most L as well. Note that all the jobs in J_i can be scheduled on the machine $M(i)$ in unit time. We obtain a scheduling σ^0 by assigning the jobs J_i to the machines M_i in the time slot used in σ to schedule the job j_i . This scheduling of I^0 is indeed a valid scheduling, and has a makespan of at most

Soundness We prove the soundness part in the lemma below.

Lemma 192. Suppose that there is a scheduling σ with makespan L . Then, we will show that there is a scheduling σ^0 with makespan at most L .

Proof. Note that there is a trivial scheduling σ where we schedule jobs in a topological sort one by one, with makespan equal to n . Thus, henceforth, we assume that $n > L$.

Let $\epsilon = \frac{1}{10n^2}$. We claim that for every $i \in [n]$, at most $2^{(m - M(i))}$ jobs in J_i are processed by machines that do not belong to M_i in the scheduling σ^0 . The proof of this claim follows from Lemma 1 of [BN15], and we present it here for the sake of completeness. Fix an index $i \in [n]$, and for ease of notation, let $M = M(i)$. First, as each job j_i has length ϵ^{-1} , and the

processing speed of each machine $M_j, j < i$ is at most $\frac{1}{i-1}$, no job in J_i is scheduled on machines $M_j, j < i$, as the makespan of J_i is at most $L < \frac{1}{i-1}$. Now, consider an integer $j \in [m], j > i$. There are $\binom{m}{j}$ machines M_j , and they have a processing speed of $\frac{1}{j}$. Thus, in time L , they can process at most

$$\frac{n \binom{m}{j}}{j} < \frac{n}{j-1} < \frac{n}{2^{(m-i)}}$$

jobs of J_i . Taking union over all $j > i$, we get that at most

$$\frac{nm}{2^{(m-i)}} < \frac{1}{10n^2} < \frac{1}{10n^2} < \frac{1}{10n^2}$$

jobs in J_i are processed by machines outside $M(i)$. In other words, for every job j of J_i , at most $\frac{1}{10n^2}$ fraction of the jobs in J_i are processed by machines outside $M(i)$.

Now, consider a scheduling of the jobs in J_i where for every $j \in [n]$, we get rid of the jobs in J_i that are processed by machines outside $M(i)$. After removing the jobs processed by other machines, we still have that for every $j \in [n]$, at least $\frac{9}{10}$ fraction of the jobs in J_i are processed. Also observe that since we are only deleting some jobs, the makespan of the new scheduling is at most L as well. Recall that processing each job j takes unit time on the machines $M(i)$.

Using this observation, we obtain a fractional scheduling of time L as follows. For every $j \in [n]$ and $t \in [L]$, define the variable $x_{j,t}$ to be the fraction of the jobs of J_i that are scheduled by the machine $M(i)$ in the time slot t . By the above discussion, we get the following properties of this fractional scheduling.

1. Every job $j \in [n]$ is almost fully processed. For every $j \in [n]$, we have

$$\sum_{t=1}^L x_{j,t} = 1$$

2. Every machine is used only for processing a single unit of job in a time slot.

$$\sum_{j \in J(i)} x_{j,t} = 1 \quad \forall i \in [m]; t \in [L]$$

3. If there is a precedence constraint $j_1 \prec j_2$ in J , j_2 's processing is done only in the time slots after j_1 is fully processed. More formally,

$$x_{j_1;t} > 0 \implies x_{j_2;t} = 0 \quad \forall t < L$$

We will now show that the fractional scheduling implies that the instance has an integral scheduling with makespan at most $2L$, thereby proving the Lemma. We will prove this in two steps: first, we modify the fractional scheduling to obtain another fractional scheduling with better structure, and then next, we use this to obtain the integral scheduling.

For a job $j \in [n]$, define the starting time e_j^s and the end time e_j^e as the minimum and the maximum times at which j is being processed.

$$e_j^s = \min \{ t : x_{j,t} > 0 \}; e_j^e = \max \{ t : x_{j,t} > 0 \}$$

Note that if we have $j_1, t_1^s > t_1^e$. We now modify the fractional scheduling to ensure that each machine processes the job with the lowest ending time first, from the available set of the jobs. More formally, for a machine m , consider the pair of jobs $j_1, j_2 \in J$ and time slot $t \in [L]$ satisfying the following conditions.

(C1) The job j_1 has lower ending time $t_1^e < t_2^e$, or $t_1^e = t_2^e$ and $l_1 < l_2$.

(C2) The job j_1 can be processed on the time slot t but the job j_2 is processed instead of finishing the job j_1 :

$$t_1^s < t < t_1^e; x_{l_2;t} > 0$$

If there are jobs j_1, j_2 and time slot t satisfying these conditions, we swap the processing times, and process the job j_1 in the time slot t instead of j_2 . More formally, let $t^0 > t$ be such that $x_{l_1;t^0} > 0$. Let $y = \min(x_{l_1;t^0}, x_{l_2;t})$. We obtain a new fractional scheduling by setting

$$x_{l_1;t^0} = x_{l_1;t^0} - y; x_{l_1;t} = x_{l_1;t} + y$$

$$x_{l_2;t^0} = x_{l_2;t^0} + y; x_{l_2;t} = x_{l_2;t} - y$$

Note that the operation does not increase the ending time of either job and does not decrease the starting time of either job and thus, results in a valid fractional scheduling respecting the precedence conditions. We repeat the swapping operations until there is no triple left where both (C1) and (C2) are true. We also update the starting and ending times of the jobs appropriately when we apply the swapping operations.

Next, we apply another transformation to the fractional scheduling by filling the empty slots in the machines, if there are any. More formally, consider a time slot $t \in [L]$ and job $j \in J$ such that the following hold.

(D1) The time slot t is not fully utilized:

$$\sum_{j \in J} x_{l_j;t} < 1$$

(D2) The job j can be scheduled on the time slot t instead of leaving the machine idle: $t < t_j^e$.

If there is a time slot t and job j such that the above two conditions hold, we fill the empty slot in the time slot t by processing the job j . Let $t^0 > t$ be such that $x_{l_j;t^0} > 0$. Let $y = \min(x_{l_j;t^0}, 1 - \sum_{j \in J} x_{l_j;t})$. We set

$$x_{l_j;t} = x_{l_j;t} + y; x_{l_j;t^0} = x_{l_j;t^0} - y$$

We repeat these operations iteratively until no empty slots can be filled. Similar to the previous case, we update the starting and ending times of the jobs appropriately.

After the two types of operations, we obtain a fractional scheduling with the following property: at every time slot, for a machine m , let $S_{i;t}$ be the set of jobs that can be scheduled in the time slot:

$$S_{i;t} := \{j \in J : t_j^s < t < t_j^e\}$$

We sort the jobs in $S_{i,t}$ as l_1, l_2, \dots, l_k by increasing order of ending times, and breaking ties based on the index. The fractional scheduling greedily schedules the jobs l_1, l_2, \dots, l_k in that order. More formally, we have

$$x_{l_1,t} = \begin{cases} 1 & \text{if } t=1 \\ x_{l_1,t}^0 & \text{otherwise} \end{cases}$$

and

$$x_{l_2,t} = \min \left\{ 1 - x_{l_1,t}, \begin{cases} 1 & \text{if } t=1 \\ x_{l_2,t}^0 & \text{otherwise} \end{cases} \right\}$$

and so on.

Our goal is to show that in this fractional scheduling that we obtained, each machine schedules at most two jobs in any time slot. In order to prove this, we first define the following parameter $P_{i,t}$, the amount of job partially completed in the machine i by the time t .

$$P_{i,t} = \sum_{l \in J(i): t_l^e > t} x_{l,t}^0$$

We claim that for every $i \in [m]; t \in [L]$, we have $P_{i,t} \leq 1$. Fix a machine $i \in [m]$. We will prove the claim by induction on t .

1. Base case when $t = 1$. If no job is processed by the machine in the time slot $t = 1$, the claim is trivially satisfied. Else, let l_1 be the job in $J(i)$ with the lowest ending time, breaking ties by the lowest index. Note that the fractional scheduling fully schedules the job l_1 in the time slot $t = 1$. As each job is processed for at least $\frac{1}{10n}$ duration, we get that $P_{i,1}$ is at most 1 .
2. Inductive proof. Suppose that the claim holds for all t and consider the time slot $t+1$. For ease of notation, let $S = S_{i,t+1}$ be the set of jobs that can be processed on the machine in the time slot $t+1$. If S is empty, the inductive claim trivially holds. Else, let $l_2 \in S$ be the job with the lowest ending time (breaking ties by the least index). Note that the modified fractional scheduling finishes the job l_2 in the time slot $t+1$. Let $x_{l_2,t}^0$ denote the amount of the job l_2 that is processed by time t , i.e., $x_{l_2,t}^0 = \sum_{s=1}^t x_{l_2,s}$. The amount of jobs that are partially finished at the end of time slot $t+1$ is at most

$$P_{i,t+1} = P_{i,t} + x_{l_2,t}^0 + (1 - x_{l_2,t+1})$$

We will now show that every machine processes at most two jobs in a time slot. Consider a machine $i \in [m]$ and time slot $t \in [L]$. Let $S_{i,t} := \{l_1, l_2, \dots, l_k\}$. By the previous claim, we know that at most $(t-1)$ portion of the job l_u is finished before time t , for every $u \in [k]$. Note that $(t-1) \leq L - \frac{1}{10n}$. Thus, the greedy fractional scheduling can only schedule at most two jobs, as each of them takes at least $\frac{1}{10n}$ time. Finally, using this observation, we can duplicate every time slot to obtain an integral scheduling with makespan at most $2L$. \square

Parameter analysis. The number of machines in the related machines scheduling instance is $M = O(m) = n^{O(m)}$, while the hardness gap is $(\log n)^{1-\epsilon}$ for every $\epsilon > 0$. By setting ϵ appropriately, we get a hardness gap of $(\log M)^{1-\epsilon}$ for the scheduling related machines with precedences problem.

10.4.1 Hypothesis of [BN15] implies superconstant hardness of the UMPS problem with unit lengths

Bazzi and Norouzi-Fard [BN15] introduced the following hypothesis and proved that it implies a superconstant hardness for scheduling related machines with precedences.

Hypothesis 193([BN15]). For every $\epsilon > 0$ and constant integers $k, Q > 0$, the following problem is NP-hard. Given a k -partite graph $G = (V_1; V_2; \dots; V_k; E_1; E_2; \dots; E_{k-1})$ with $|V_i| = n$ for all $1 \leq i \leq k$, and E_i being the set of edges between V_i and V_{i+1} for every $1 \leq i < k$, distinguish between the two cases:

1. (YES case) Every V_i can be partitioned into $V_{i,0}; V_{i,1}; \dots; V_{i,Q-1}$ such that
 - There is no edge between V_{i,j_1} and V_{i+1,j_2} for all $1 \leq i < k; j_1 > j_2 \in [Q]$.
 - $|V_{i,j}| \leq \frac{1}{Q}n$ for all $i \in [k]; j \in [Q]$.
2. (NO case) For every $1 < i \leq k$, and any two sets S, T with $S \subseteq V_i, T \subseteq V_{i-1}, |S| = |T| = n$, there is an edge between S and T .

We now prove that the above hypothesis implies that it is NP-hard to obtain a constant factor approximation algorithm for the UMPS problem, even when all the jobs have unit length.

Reduction. Given an instance of k -partite problem G , we output an instance of the UMPS problem as follows: there are n unit sized jobs in G , one job corresponding to each vertex of G . There are k machines, and all the jobs $v, i \in [k]$ can only be scheduled on the machine i . For every edge $e = (u; v)$ in the graph such that $u \in V_i; v \in V_{i+1}$, we have a precedence condition $u \prec v$ in G . We choose the parameter $Q = k$, and $\epsilon = \frac{1}{k}$.

Completeness. Suppose that the YES case of Hypothesis 193 holds i.e., there is a partition of V_i into $V_{i,0}; V_{i,1}; \dots; V_{i,Q-1}$ respecting the two conditions above. Then, we claim that there is a scheduling of G with makespan at most n . For every machine $i \in [k]$, we schedule the jobs in $V_{i,0}$ (in arbitrary order), and then the jobs $V_{i,1}$ (in arbitrary order) and so on. However, we start the execution of the jobs in $V_{i,0}$ at time t_i , and then, execute the jobs $V_{i,1}$ immediately after the execution of the jobs in $V_{i,0}$ for all $i \in [k]$. The parameters $t_i, i \in [k]$ are chosen such that for every pair of jobs $u; v$ with $u \prec v$, u is guaranteed to have scheduled before v . In particular, we choose $t_i = (i-1)n + \frac{1}{Q}$.

We now prove that this results in a valid scheduling that respects the precedence conditions. Consider a pair of jobs $u; v$ such that $u \in V_i; v \in V_{i+1}$ such that $u \prec v$. As the k -partite graph satisfies the YES condition, we have integers j_1, j_2 such that $u \in V_{i,j_1}$ and $v \in V_{i+1,j_2}$, and $j_1 > j_2$. Note that u is processed by time at most

$$t_u = t_i + |V_{i,0}| + |V_{i,1}| + \dots + |V_{i,j_1}|$$

Furthermore is processed only after time

$$t_v = t_{i+1} + jV_{i+1;0} + jV_{i+1;1} + \dots + jV_{i+1;j_2-1}$$

We have

$$\begin{aligned} t_v - t_u &= t_{i+1} + jV_{i+1;0} + jV_{i+1;1} + \dots + jV_{i+1;j_2-1} - (t_i + jV_{i;0} + jV_{i;1} + \dots + jV_{i;j_1-1}) \\ &= n + \frac{1}{Q} + jV_{i+1;0} + jV_{i+1;1} + \dots + jV_{i+1;j_2-1} - (n + jV_{i;j_1-1} + jV_{i;j_1+2} + \dots + jV_{i;j_1-1}) \\ &= n + \frac{1}{Q} + j_2 \frac{(1-n)}{Q} - n \frac{(Q-j_1-1)(1-n)}{Q} \\ &= n + \frac{1}{Q} + j_1 \frac{(1-n)}{Q} - j_1 \frac{(1-n)}{Q} + n + \frac{(1-n)}{Q} = 0 \end{aligned}$$

Thus, the schedule is a valid scheduling. The makespan of this scheduling is at most $t_k + n = (k-1)n + \frac{1}{Q} + n \leq 3n$.

Soundness Suppose that the NO case of Hypothesis 193 holds. We claim that in this case, the makespan of σ is at least $(1-\frac{2}{Q})kn$. For every $i \in [k]$, let s_i denote the time at which the machine i has finished $(1-\frac{1}{Q})n$ jobs of V_i . For an index $i \in [k]$, let $S(i) \subseteq V_i$ denote the set of jobs that are not processed by the time s_i . By the definition of s_i , we have $|S(i)| \leq \frac{n}{Q}$. By the NO case of Hypothesis 193, we get that there are at least $(1-\frac{1}{Q})n$ jobs in V_{i+1} that have dependencies in $S(i)$. Note that all these jobs can be scheduled only after

$$s_{i+1} = s_i + (1-\frac{2}{Q})n \quad \forall i \in [k-1]$$

Summing over all, we get that the makespan of the scheduling is at least $(1-\frac{2}{Q})kn$, which is at least $\frac{kn}{2}$ when $k \geq 4$. By choosing k large enough, this completes the proof that assuming Hypothesis 193, it is NP-hard to obtain a $(1-\frac{1}{Q})$ factor approximation algorithm for the UMPS problem when the jobs have unit lengths.

Chapter 11

Conclusion

We conclude by mentioning a few directions for further research.

Boolean PCSP Dichotomy Can we show that every Boolean PCSP is either in P or is NP-Hard? Problem 194. When can a Boolean PCSP be solved in polynomial time? Is it true that every Boolean PCSP can be solved in polynomial time or is NP-Hard?

For the case of CSPs (on general domains), a CSP has a polynomial time algorithm if and only if it has a cyclic polymorphism of arity at least 3. The hardness part was proved in the early 2000s itself [BJK05] and proving the algorithmic part was the main challenge [Bul17; Zhu20]. However, for Promise CSPs, even in the Boolean case, we do not have a candidate characterization of polymorphisms that leads to polynomial time algorithms.

On the algorithmic side, the best result is that symmetric polymorphisms of arbitrarily large arity lead to algorithms [Bra+20] using a combination of the basic LP relaxation and the affine relaxation. Could it be true that using the basic SDP relaxation together with the affine relaxation gives an optimal algorithm for all Boolean PCSPs? We shed some light on this question in Chapter 6 where we study the power of the basic SDP relaxation for PCSPs. But our result applies to the basic SDP alone, and a potential avenue to understand the power of basic SDP together with the Affine relaxation to study the minimum M_{SDP+A} such that a PCSP can be solved by the basic SDP relaxation together with the affine relaxation if and only if there is a minion homomorphism from M_{SDP+A} to $\text{Pol}(\Gamma)$.

On the hardness side, as seen in Chapter 3, the rich conjecture [BKM21] could be of help in obtaining improved NP-Hardness results. In Chapter 3, we showed that when the Boolean PCSP contains the predicate γ , i.e., when the polymorphisms are all monotone functions, if every polymorphism contains a coordinate of high Shapley influence, then the underlying PCSP is NP-Hard (under the rich-to-1 conjecture). Can we extend this result to arbitrary Boolean functions using a suitable generalization of Shapley influence?

NP-Hardness of Approximate Graph Coloring. Despite great progress on Promise CSPs,

¹A function f of arity $l \geq 2$ is said to be cyclic iff $f(x_1; x_2; \dots; x_l) = f(x_2; x_3; \dots; x_l; x_1) = \dots = f(x_l; x_1; \dots; x_{l-1})$ for every $x_1; x_2; \dots; x_l$.

²A Minion homomorphism $f: M_1 \rightarrow M_2$ (formally defined in Chapter 6) is a mapping that preserves the arity of the elements and also commutes with taking minors.

we still do not know if it is NP-Hard to ϵ -color a k -colorable graph in polynomial time. More generally, let $(c; s)$ -approximate graph coloring be the computational problem of coloring a graph that is promised to be k -colorable with s colors.

Problem 195. Prove that $(3; s)$ -approximate graph coloring is NP-Hard for every constant s .

In Chapter 4, we have proved that the $(k-1)$ conjecture [Kho02b] for any constant k implies that it is NP-Hard to color k -colorable graphs with $O(1)$ colors, thus resolving Problem 195. The imperfect completeness version of the $(k-1)$ conjecture when $k = 2$ was proved recently in a breakthrough series of works [KMS17; Din+18b; Din+18a; KMS18]. However, their result starts with the hardness of linear equations as a starting point, and thus, does not extend to the case when the problem has perfect completeness. Can we obtain a proof (perhaps with larger k) conjecture by using a different problem as a starting point, and studying analogous objects to Grassman graphs of [KMS18]?

A different approach to Problem 195 is using the PCSP machinery developed recently. Barto, Bulin, Krokhin, Opsal [Bar+21] proved that k -coloring a k -colorable graph is NP-Hard. They achieve this by showing that there is a minion homomorphism from the polymorphisms of this PCSP to the polymorphisms of $(k-1)$ -coloring a k -uniform hypergraphs. By the result of Dinur, Regev, Smyth [DRS05], it is NP-Hard to color a k -colorable k -uniform hypergraph with $O(1)$ colors, thus implying the hardness of $(k, 5)$ -approximate graph coloring. Approximate graph coloring is a special case of a more general graph homomorphism problem where the input is a pair of graphs H_1, H_2 such that there is a homomorphism $H_1 \rightarrow H_2$, given an input graph G with the promise that $G \rightarrow H_1$, can we find a homomorphism $G \rightarrow H_2$? $(c; s)$ -approximate graph coloring is the case when $H_1 = K_c; H_2 = K_s$. Recently, Krokhin, Opsal, Wrochna, Zivny [Kro+20] proved the NP-Hardness of the graph homomorphism problem when $H_1 = K_3; H_2 = K_2$ for every odd integer $c \geq 3$. A possible avenue towards proving Problem 195 is to generalize their techniques for the case when H_2 is a larger clique.

Robust algorithms for PCSPs. In Chapter 6, we initiated the study of robust algorithms for PCSPs. Characterizing which PCSPs is interesting on its own, but more so in connection with the power of SDP algorithms for PCSPs (Conjecture 65).

The main goal is to answer the following question:

Problem 196. Which Promise CSPs have Robust Algorithms? Is it the same class as Promise CSPs that can be solved by the basic SDP relaxation?

As we proved in Chapter 6, the existence of robust algorithms for a PCSP is characterized by its polymorphisms. We have shown that having ATG or MAJ polymorphisms of all odd arities leads to algorithms. On the other hand, the $\text{CSP}_{\text{Parity}}$, which has Parity polymorphisms of all odd arities, does not admit a robust algorithm. This begs the question: which polymorphism families lead to robust algorithms? Intuitively, the polymorphisms should be robust to noise, and the challenge is to precisely characterize the notion of noise stability that leads to algorithms and incorporate it into the robust algorithms.

On the hardness side, proving integrality gaps for the basic SDP relaxation and using Raghavendra's [Rag08] framework is a general technique to show robust hardness. Using Lemma 86, we can reduce the problem of finding integrality gaps for the basic SDP of a PCSP to finding sphere

colorings $f : S^n \rightarrow [r]$ satisfying certain structural properties. We have used results from sphere Ramsey theory to answer such questions, leading to robust hardness for some classes of PCSPs, but the general problem of finding specific structures in sphere colorings is wide open.

Inapproximability of Related Machines Scheduling with Precedences
 In the related machine scheduling, there are n jobs with processing times p_1, p_2, \dots, p_n with precedence constraints between them. That is, we are given a DAG over these jobs and if there is an edge from job j_1 to j_2 , the processing of j_2 can only begin after j_1 finishes. There are m machines, each with speed $s_i, i \in [m]$. A job j on a machine i takes time $\frac{p_j}{s_i}$. The objective is to schedule the jobs on the machines to minimize the makespan. The problem admits a factor approximation algorithm [CB01; CS99]. On the hardness side, the best known inapproximability is a factor 2 [Sve10; BK09], assuming a variant of the Unique Games Conjecture. Whether we can get a better approximation algorithm for the problem remained a long-standing open problem in scheduling theory, and has been asked in multiple influential surveys [SW99a; Ban17; WS11].

Problem 197. Is there an $O(1)$ factor approximation algorithm for related machine scheduling with precedences problem?

As we proved in Chapter 10, improved hardness of the Unique Machine Precedence Scheduling (UMPS) problem Conjecture 185 implies poly logarithmic hardness of scheduling related machines with precedences problem. A potential avenue to showing the hardness of the UMPS problem is via using the rich-to-1 conjecture [BKM21] in the framework of Bansal and Khot [BK09].

Bibliography

- [ABP19] Per Austrin, Amey Bhangale, and Aditya Potukuchi. “Simplified inapproximability of hypergraph coloring via t-agreeing families”. CoRRabs/1904.01163 (2019). arXiv: 1904.01163 .
- [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, pp. 1479–1495.
- [AD22] Albert Atserias and Víctor Dalmau. “Promise Constraint Satisfaction and Width”. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) SIAM. 2022, pp. 1129–1153.
- [AGH17] Per Austrin, Venkatesan Guruswami, and Johan Hastad. “(2+)-SAT Is NP-hard”. In: SIAM J. Comput. 46.5 (2017), pp. 1554–1573.
- [Aha01] Ron Aharoni. “Ryser's Conjecture for Tripartite 3-Graphs”. Comb. 21.1 (2001), pp. 1–4.
- [AHK96] Ron Aharoni, Ron Holzman, and Michael Krivelevich. “On a Theorem of Alon Covers in tau-Partite Hypergraphs”. Combinatorica 16.2 (1996), pp. 149–174.
- [AKS11] Per Austrin, Subhash Khot, and Muli Safra. “Inapproximability of Vertex Cover and Independent Set in Bounded Degree Graphs”. Theory Comput. 7.1 (2011), pp. 27–43.
- [Alo+98] Noga Alon, Yossi Azar, Avrim Blum, Mihalis Christofides, Mihalis Yannakakis, and Gerhard J. Woeginger. “On-Line and Off-Line Approximation Algorithms for Vector Covering Problems”. Algorithmica 21.1 (1998), pp. 104–118.
- [APV16] Andris Ambainis, Krišjānis Pūsis, and Jevgēnijs Vihrovs. “Sensitivity Versus Certificate Complexity of Boolean Functions”. In: Proceedings of the 11th International Computer Science Symposium on Computer Science — Theory and Applications - Volume 969. ICSR 2016. St. Petersburg, Russia, 2016, pp. 16–28.
- [Aro+98] Sanjeev Arora, Carsten Lund, Rajeev Motwani, Madhu Sudan, and Mario Szegedy. “Proof Verification and the Hardness of Approximation Problems”. In: J. ACM 45.3 (1998), pp. 501–555.
- [AS98] Sanjeev Arora and Shmuel Safra. “Probabilistic Checking of Proofs: A New Characterization of NP”. In: Journal of the ACM 45.1 (1998), pp. 70–122.

- [AZ20] Ron Aharoni and Shira Zerbib. “A generalization of Tuza’s conjecture” *Journal of Graph Theory* 94.3 (2020), pp. 445–462.
- [Aza+13] 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 2013*, pp. 961–970.
- [Ban17] Nikhil Bansal. “Scheduling Open Problems: Old And New.” *MAPSP 2017* www.mapsp2017.ma.tum.de/MAPSP2017-Bansal.pdf (2017).
- [Bar+21] Libor Barto, Jakub Bůh, Andrei A. Krokhin, and Jakub Opral. “Algebraic Approach to Promise Constraint Satisfaction”. *JnACM* 68.4 (2021), 28:1–28:66.
- [Bar18a] Libor Barto. Personal communication. 2018.
- [Bar18b] Libor Barto. “Cyclic operations in promise constraint satisfaction problems”. In: *Dagstuhl workshop The Constraint Satisfaction Problem: Complexity and Approximability, Schloss Dagstuhl, Germany* (2018). Available at https://www2.karlin.mff.cuni.cz/~barto/Articles/Barto_Dagstuhl18.pdf.
- [BCD20] Patrick Bennett, Ryan Cushman, and Andrzej Dudek. “Closing the Random Graph Gap in Tuza’s Conjecture Through the Online Triangle Packing Process”. *arXiv: 2007.04478 [math.CO]* .
- [BCS09] Nikhil Bansal, Alberto Caprara, and Maxim Sviridenko. “A New Approximation Method for Set Covering Problems, with Applications to Multidimensional Bin Packing”. In: *SIAM J. Comput* 39.4 (2009), pp. 1256–1278.
- [BE51] NG de Bruijn and P Erdos. “A colour problem for finite graphs and a problem in the theory of relations”. In: *Indagationes Mathematicae* 3 (1951), pp. 371–373.
- [BEK16] Nikhil Bansal, Marek Elías, and Arindam Khan. “Improved Approximation for Vector Bin Packing”. In: *Proceedings of the Twenty-Seventh Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2016*, pp. 1561–1579.
- [BFG19] Fábio Botler, Cristina G. Fernandes, and Juan Cruz. “On Tuza’s Conjecture for Triangulations and Graphs with Small Treewidth”. *Proceedings of the tenth Latin and American Algorithms, Graphs and Optimization Symposium, LAGOS 2019, Belo Horizonte, Brazil, June 2-7, 2019*. Vol. 346. *Electronic Notes in Theoretical Computer Science*. Elsevier, 2019, pp. 171–183.
- [BFM10] Tom Bohman, Alan M. Frieze, and Dhruv Mubayi. “Coloring-free hypergraphs”. In: *Random Struct. Algorithms* 36.1 (2010), pp. 11–25.
- [BG16] Joshua Brakensiek and Venkatesan Guruswami. “New Hardness Results for Graph and Hypergraph Colorings”. In: *31st Conference on Computational Complexity, CCC 2016*, 14:1–14:27.
- [BG17] Joshua Brakensiek and Venkatesan Guruswami. “The Quest for Strong Inapproximability Results with Perfect Completeness”. *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2017, August 16-18, 2017, Berkeley, CA, USA* 2017, 4:1–4:20.

- [BG21a] Joshua Brakensiek and Venkatesan Guruswami. “Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy”. *SIAM J. Comput.* 50.6 (2021), pp. 1663–1700.
- [BG21b] Joshua Brakensiek and Venkatesan Guruswami. “Promise Constraint Satisfaction: Algebraic Structure and a Symmetric Boolean Dichotomy”. *SIAM J. Comput.* 50.6 (2021), pp. 1663–1700.
- [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*, pp. 152–174.
- [BGS21] Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. “Conditional Dichotomy of Boolean Ordered Promise CSPs”. *46th International Colloquium on Automata, Languages, and Programming, ICALP 2021, July 12-16, 2021, Glasgow, Scotland (Virtual Conference)*. Vol. 198. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 37:1–37:15.
- [BGS98] Mihir Bellare, Oded Goldreich, and Madhu Sudan. “Free Bits, PCPs, and Nonapproximability-Towards Tight Results”. In: *SIAM J. Comput.* 27.3 (1998), pp. 804–915.
- [Bha18] Amey Bhangale. “NP-Hardness of Coloring 2-Colorable Hypergraph with Poly-Logarithmically Many Colors”. In: *45th International Colloquium on Automata, Languages, and Programming* 2018, 15:1–15:11.
- [BJK05] Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. “Classifying the Complexity of Constraints Using Finite Algebras”. In: *SIAM J. Comput.* 34.3 (2005), pp. 720–742.
- [BK09] Nikhil Bansal and Subhash Khot. “Optimal Long Code Test with One Free Bit”. In: *50th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2009, October 25-27, 2009, Atlanta, Georgia, USA*. IEEE Computer Society, 2009, pp. 453–462.
- [BK10] Nikhil Bansal and Subhash Khot. “Inapproximability of Hypergraph Vertex Cover and Applications to Scheduling Problems”. In: *Automata, Languages and Programming, 37th International Colloquium, ICALP 2010, Bordeaux, France, July 6-10, 2010, Proceedings, Part 2* 2010, pp. 250–261.
- [BK14a] 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*, pp. 13–25.
- [BK14b] Libor Barto and Marcin Kozik. “Constraint Satisfaction Problems Solvable by Local Consistency Methods”. In: *ACM* 61.1 (2014), 3:1–3:19.
- [BK16] Libor Barto and Marcin Kozik. “Robustly Solvable Constraint Satisfaction Problems”. In: *SIAM J. Comput.* 45.4 (2016), pp. 1646–1669.
- [BKM21] Mark Braverman, Subhash Khot, and Dor Minzer. “On Rich 2-to-1 Games”. In: *12th Innovations in Theoretical Computer Science Conference, ITCS 2021*. LIPIcs. 2021, 27:1–27:20.

- [BKW17] Libor Barto, Andrei Krokhin, and Ross Willard. “Polymorphisms, and How to Use Them”. In: *The Constraint Satisfaction Problem: Complexity and Approximability*. Vol. 7. Dagstuhl Follow-Ups. Dagstuhl, Germany: Schloss Dagstuhl–Leibniz-Zentrum fuer Informatik, 2017, pp. 1–44.
- [BN15] Abbas Bazzi and Ashkan Norouzi-Fard. “Towards Tight Lower Bounds for Scheduling Problems”. In: *Algorithms - ESA 2015 - 23rd Annual European Symposium*, Patras, Greece, September 14-16, 2015, *Proceedings* 9294. Lecture Notes in Computer Science. Springer, 2015, pp. 118–129.
- [Bra+20] Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Zivny. “The Power of the Combined Basic Linear Programming and Affine Relaxation for Promise Constraint Satisfaction Problems”. *SIAM J. Comput.* 49.6 (2020), pp. 1232–1248.
- [Bra+21] Joshua Brakensiek, Neng Huang, Aaron Potechin, and Uri Zwick. “On the mysteries of MAX NAE-SAT”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)* SIAM. 2021, pp. 484–503.
- [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*, pp. 196–203.
- [Bul17] Andrei A. Bulatov. “A Dichotomy Theorem for Nonuniform CSPs”. *58th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2017* Computer Society, 2017, pp. 319–330.
- [BZ21] Alex Brandts and Stanislav Zivny. “Beyond PCSP(1-in-3, NAE)”. In: *48th International Colloquium on Automata, Languages, and Programming, ICALP 2021*, July 12-16, 2021, Glasgow, Scotland (Virtual Conference). 198. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2021, 121:1–121:14.
- [Cae94] D. de Caen. “The current status of Tait’s problem on hypergraphs”. *Extremal Problems for Finite Sets* 1991, Bolyai Soc. Math. Stud., Vol. 3, pp. 187–197, eds Bolyai Math. Soc., Budapest (1994).
- [Car+12] Jean Cardinal, Marek Karpinski, Richard Schmied, and Claus Viehmann. “Approximating vertex cover in dense hypergraphs”. *Journal of discrete algorithms* 13 (2012), pp. 67–77.
- [CB01] Chandra Chekuri and Michael A. Bender. “An Efficient Approximation Algorithm for Minimizing Makespan on Uniformly Related Machines”. In: *Algorithms* 41.2 (2001), pp. 212–224.
- [CC06] Miroslav Chlebik and Janka Chlebikova. “Inapproximability Results for Orthogonal Rectangle Packing Problems with Rotations”. *Algorithms and Complexity, 6th Italian Conference, CIAC 2006* 2006, pp. 199–210.
- [CD72] J Csimá and B.N Datta. “The DAD theorem for symmetric non-negative matrices”. In: *Journal of Combinatorial Theory, Series A* 12.1 (1972), pp. 147–152.
- [CFS08] L. Sunil Chandran, Mathew C. Francis, and Naveen Sivadasan. “Boxicity and maximum degree”. In: *J. Comb. Theory, Ser. B* 98.2 (2008), pp. 443–445.

- [Cha+20] Parinya Chalermsook, Samir Khuller, Pattara Sukprasert, and Sumedha Uniyal. “Multi-transversals for Triangles and the Tuza's Conjecture”. *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020*. Ed. by Shuchi Chawla. SIAM, 2020, pp. 1955–1974.
- [Cha16] Siu On Chan. “Approximation Resistance from Pairwise-Independent Subgroups”. In: *J. ACM*63.3 (2016), 27:1–27:32.
- [Cho+11] Mosharaf Chowdhury, Matei Zaharia, Justin Ma, Michael I. Jordan, and Ion Stoica. “Managing data transfers in computer clusters with orchestra”. *Proceedings of the ACM SIGCOMM 2011 Conference on Applications, Technologies, Architectures, and Protocols for Computer Communications, Toronto, ON, Canada, August 15-19, 2011*. ACM, 2011, pp. 98–109.
- [Chr+17] Henrik I. Christensen, Arindam Khan, Sebastian Pokutta, and Prasad Tetali. “Approximation and online algorithms for multidimensional bin packing: A survey”. In: *Comput. Sci. Res.*24 (2017), pp. 63–79.
- [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*. 2001, pp. 557–566.
- [CK04] Chandra Chekuri and Sanjeev Khanna. “On Multidimensional Packing Problems”. In: *SIAM J. Comput.*33.4 (2004), pp. 837–851.
- [CKL20] Vincent Cohen-Addad, Karthik C.S., and Euiwoong Lee. “On Approximability of k-means, k-median, and k-minsum Clustering”. *Manuscript*(2020).
- [CKL21] Vincent Cohen-Addad, Karthik C. S., and Euiwoong Lee. “On Approximability of Clustering Problems Without Candidate Centers”. *Proceedings of the ACM-SIAM Symposium on Discrete Algorithms*. To appear. 2021.
- [CL99] Fan Chung and Linyuan Lu. “An upper bound for the Turán number $t_3(n, 4)$ ”. In: *Journal of Combinatorial Theory, Series B*77.2 (1999), pp. 381–389.
- [CMM09] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. “Near-optimal algorithms for maximum constraint satisfaction problems”. *ACM Trans. Algorithms* 5.3 (2009), 32:1–32:14.
- [CS00] Artur Czumaj and Christian Scheideler. “A new algorithm approach to the general Lovász local lemma with applications to scheduling and satisfiability problems (extended abstract)”. In: *Proceedings of the Thirty-Second Annual ACM Symposium on Theory of Computing, May 21-23, 2000, Portland, OR, USA*. ACM, 2000, pp. 38–47.
- [CS99] Fabřan A. Chudak and David B. Shmoys. “Approximation Algorithms for Precedence-Constrained Scheduling Problems on Parallel Machines that Run at Different Speeds”. In: *J. Algorithms*30.2 (1999), pp. 323–343.
- [Cyg13] Marek Cygan. “Improved Approximation for 3-Dimensional Matching via Bounded Pathwidth Local Search”. In: *54th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2013, Berkeley, CA, USA*. 2013.

- [CZ22a] Lorenzo Ciardo and Stanislav Žitný. “CLAP: A New Algorithm for Promise CSPs”. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) SIAM. 2022, pp. 1057–1068.
- [CZ22b] Lorenzo Ciardo and Stanislav Žitný. “The Sherali-Adams Hierarchy for Promise CSPs through Tensors”. In: CoRRabs/2203.02478 (2022). arXiv:2203.02478 .
- [Dav+20] Sami Davies, Janardhan Kulkarni, Thomas Rothvoss, Jakub Tarnawski, and Yihao Zhang. “Scheduling with Communication Delays via LP Hierarchies and Clustering”. In: To appear at 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2020), 16-19 November 2020, Durham, North Carolina, USA, Proceedings (2020).
- [Dav+21] Sami Davies, Janardhan Kulkarni, Thomas Rothvoss, Jakub Tarnawski, and Yihao Zhang. “Scheduling with Communication Delays via LP Hierarchies and Clustering II: Weighted Completion Times on Related Machines”. In: Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms, SODA 2021, Virtual Conference, January 10 - 13, 2021. Ed. by Daniel Marx. SIAM, 2021, pp. 2958–2977.
- [Dav+22] Sami Davies, Janardhan Kulkarni, Thomas Rothvoss, Sai Sandeep, Jakub Tarnawski, and Yihao Zhang. “On the Hardness of Scheduling With Non-Uniform Communication Delays”. In: Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA) SIAM. 2022, pp. 316–328.
- [DDS17] Anindya De, Ilias Diakonikolas, and Rocco A. Servedio. “The Inverse Shapley value problem”. In: Games Econ. Behav. 105 (2017), pp. 122–147.
- [Din+05] Irit Dinur, Venkatesan Guruswami, Subhash Khot, and Oded Regev. “A New Multi-layered PCP and the Hardness of Hypergraph Vertex Cover”. In: SIAM J. Comput. 34.5 (2005), pp. 1129–1146.
- [Din+18a] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. “On non-optimally expanding sets in Grassmann graphs”. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, 2018, pp. 940–951.
- [Din+18b] Irit Dinur, Subhash Khot, Guy Kindler, Dor Minzer, and Muli Safra. “Towards a proof of the 2-to-1 games conjecture?”. In: Proceedings of the 50th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2018, 2018, pp. 376–389.
- [Din07] Irit Dinur. “The PCP theorem by gap amplification”. In: Journal of the ACM 54.3 (2007), p. 12.
- [DLR95] Richard A. Duke, Hanno Lefmann, and Vojtěch Rödl. “On Uncrowded Hypergraphs”. In: Random Struct. Algorithms 6.2/3 (1995), pp. 209–212.
- [DMR09] Irit Dinur, Elchanan Mossel, and Oded Regev. “Conditional Hardness for Approximate Coloring”. In: SIAM J. Comput. 39.3 (2009), pp. 843–873.
- [DR06] Irit Dinur and Omer Reingold. “Assignment Testers: Towards a Combinatorial Proof of the PCP Theorem”. In: SIAM J. Comput. 36.4 (2006), pp. 975–1024.

- [DRS05] Irit Dinur, Oded Regev, and Clifford D. Smyth. “The Hardness of 3-Uniform Hypergraph Coloring”. In: *Combinatorica* 25.5 (2005), pp. 519–535.
- [DS05] Irit Dinur and Samuel Safra. “On the Hardness of Approximating Minimum Vertex Cover”. In: *Annals of Mathematics* 162.1 (2005), pp. 439–485.
- [DS14] Irit Dinur and David Steurer. “Analytical approach to parallel repetition”. *Symposium on Theory of Computing, STOC 2014*, pp. 624–633.
- [Erd81] Paul Erdős. “On the combinatorial problems which I would most like to see solved”. In: *Combinatorica* 1.1 (1981), pp. 25–42.
- [Erd88] Paul Erdős. “Problems and Results in Combinatorial Analysis and Graph Theory”. In: *Graph Theory and Applications, Proceedings of the First Japan Conference on Graph Theory and Application* (1988), pp. 81–92.
- [Fei98] Uriel Feige. “A Threshold of $\ln 2$ for Approximating Set Cover”. In: *J. ACM* 45.4 (1998), pp. 634–652.
- [Fic+19] Miron Ficak, Marcin Kozik, Miroslav Oštržanek, and Szymon Stankiewicz. “Dichotomy for Symmetric Boolean PCSPs”. In: *16th International Colloquium on Automata, Languages, and Programming, ICALP 2019*. LIPIcs. 2019, 57:1–57:12.
- [FK15] Alan Frieze and Micha Karoński. *Introduction to Random Graphs*. Cambridge University Press, 2015.
- [FK91] H. Furstenberg and Y. Katznelson. “A density version of the Hales-Jewett theorem.” English. In: *J. Anal. Math.* 57 (1991), pp. 64–119. ISSN: 0021-7670; 1565-8538/e.
- [FKN02] Ehud Friedgut, Gil Kalai, and Assaf Naor. “Boolean functions whose Fourier transform is concentrated on the first two levels”. In: *Adv. in Applied Math* 29 (2002), pp. 427–437.
- [FM08] Alan M. Frieze and Dhruv Mubayi. “On the Chromatic Number of Simple Triangle-Free Triple Systems”. In: *Electron. J. Comb* 15.1 (2008).
- [FM13] Alan M. Frieze and Dhruv Mubayi. “Coloring simple hypergraphs”. In: *Comb. Theory, Ser. B* 103.6 (2013), pp. 767–794.
- [Fra+17] Nevena Francetic, Sarada Herke, Brendan D. McKay, and Ian M. Wanless. “On Ryser’s conjecture for linear intersecting multipartite hypergraphs”. In: *J. Comb.* 61 (2017), pp. 91–105.
- [Fre04] Robert M Freund. “Introduction to semidefinite programming (SDP)”. Massachusetts Institute of Technology (2004), pp. 11–12. URL: https://ocw.mit.edu/courses/6-25j-introduction-to-mathematical-programming-fall-2009/08bbc2660764c4f61bde5363ae134339_MIT6_25JF09_SDP.pdf.
- [FS02] Uriel Feige and Christian Scheideler. “Improved Bounds for Acyclic Job Shop Scheduling”. In: *Comb.* 22.3 (2002), pp. 361–399.
- [FV98] Tomás Feder and Moshe Y. Vardi. “The Computational Structure of Monotone Monadic SNP and Constraint Satisfaction: A Study through Datalog and Group Theory”. In: *SIAM J. Comput.* 28.1 (1998), pp. 57–104.

- [Gar+76] M. R. Garey, Ronald L. Graham, David S. Johnson, and Andrew Chi-Chih Yao. "Resource Constrained Scheduling as Generalized Bin Packing". *Comb. Theory, Ser. A*21.3 (1976), pp. 257–298.
- [Gar18] Shashwat Garg. "Quasi-PTAS for Scheduling with Precedences using LP Hierarchies". In: 45th International Colloquium on Automata, Languages, and Programming, ICALP 2018, July 9-13, 2018, Prague, Czech Republic, 2018, pp. 59:1–59:13.
- [GCL18] Yuanxiang Gao, Li Chen, and Baochun Li. "Spotlight: Optimizing Device Placement for Training Deep Neural Networks". In: Proceedings of the 35th International Conference on Machine Learning, vol. 80. Proceedings of Machine Learning Research. Stockholmsräsan, Stockholm Sweden: PMLR, 2018, pp. 1676–1684.
- [GHS02] Venkatesan Guruswami, Johan Hastad, and Madhu Sudan. "Hardness of Approximate Hypergraph Coloring". In: *SIAM J. Comput.*31.6 (2002), pp. 1663–1686.
- [Gir+08] Rodolphe Giroudeau, Jean-Claude König, Farida Kamila Moulai, and Éric Palaysi. "Complexity and approximation for precedence constrained scheduling problems with large communication delays". In: *Theor. Comput. Sci.*401.1-3 (2008), pp. 107–119.
- [GK99] Sudipto Guha and Samir Khuller. "Greedy strikes back: Improved facility location algorithms". In: *Journal of algorithms*31.1 (1999), pp. 228–248.
- [GL17] Venkatesan Guruswami and Euiwoong Lee. "Inapproximability of Transversal/Packing". In: *SIAM J. Discret. Math.*31.3 (2017), pp. 1552–1571.
- [GL18] Venkatesan Guruswami and Euiwoong Lee. "Strong Inapproximability Results on Balanced Rainbow-Colorable Hypergraphs". *Combinatorica*38.3 (2018), pp. 547–599.
- [GO05] Venkatesan Guruswami and Ryan O'Donnell. "The PCP Theorem and Hardness of Approximation: Notes on Lectures 7-19". <https://courses.cs.washington.edu/courses/cse533/05au/>. 2005.
- [Gol+01] Leslie Ann Goldberg, Mike Paterson, Aravind Srinivasan, and Elizabeth Sweedyk. "Better Approximation Guarantees for Job-Shop Scheduling". *SIAM J. Discret. Math.*14.1 (2001), pp. 67–92.
- [GOS20] Venkatesan Guruswami, Jakub Opršal, and Sai Sandeep. "Revisiting Alphabet Reduction in Dinur's PCP". In: Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2020, August 17-19, 2020, Virtual Conference, vol. 176. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, pp. 34:1–34:14.
- [Gra+79] R. L. Graham, E. L. Lawler, J. K. Lenstra, and A. H. G. Rinnooy Kan. "Optimization and approximation in deterministic sequencing and scheduling: a survey". *Ann. Discrete Math.*4 (1979), pp. 287–326.
- [Gra66] R. L. Graham. "Bounds for Certain Multiprocessing Anomalies". *Bell System Technical Journal*45.9 (1966), pp. 1563–1581.
- [GS17] Venkatesan Guruswami and Rishi Saket. "Hardness of Rainbow Coloring Hypergraphs". In: 37th IARCS Annual Conference on Foundations of Software Technology

- and Theoretical Computer Science, FSTTCS 2017, December 11-15, 2017, Kanpur, India. 2017, 33:33–33:15.
- [GS20a] Venkatesan Guruswami and Sai Sandeep. “d-To-1 Hardness of Coloring 3-Colorable Graphs with $O(1)$ Colors”. In: 47th International Colloquium on Automata, Languages, and Programming, ICALP 2020, July 8-11, 2020, Saarbrücken, Germany (Virtual Conference) Vol. 168. LIPIcs. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2020, 62:1–62:12.
- [GS20b] Venkatesan Guruswami and Sai Sandeep. “Rainbow Coloring Hardness via Low Sensitivity Polymorphisms”. In: SIAM J. Discret. Math. 34.1 (2020), pp. 520–537.
- [GS22] Venkatesan Guruswami and Sai Sandeep. “Approximate Hypergraph Vertex Cover and generalized Tuza’s conjecture”. Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA). SIAM, 2022, pp. 927–944.
- [GSS15] Venkatesan Guruswami, Sushant Sachdeva, and Rishi Saket. “Inapproximability of Minimum Vertex Cover on k -Uniform k -Partite Hypergraphs”. SIAM J. Discret. Math. 29.1 (2015), pp. 36–58.
- [Guo+12] Zhenyu Guo, Xuepeng Fan, Rishan Chen, Jiaying Zhang, Hucheng Zhou, Sean McDirmid, Chang Liu, Wei Lin, Jingren Zhou, and Lidong Zhou. “Spotting code optimizations in data-parallel pipelines through PeriSCOPE”. Presented as part of the 10th USENIX Symposium on Operating Systems Design and Implementation (OSDI 12) 2012, pp. 121–133.
- [GZ12] Venkatesan Guruswami and Yuan Zhou. “Tight Bounds on the Approximability of Almost-Satisfiable Horn SAT and Exact Hitting Set”. In: Theory Comput. 8.1 (2012), pp. 239–267.
- [Has01] Johan Hastad. “Some optimal inapproximability results”. In: ACM 48.4 (2001), pp. 798–859.
- [Hax99] Penny E. Haxell. “Packing and covering triangles in graphs”. Discret. Math. 195.1-3 (1999), pp. 251–254.
- [HCG12] Chi-Yao Hong, Matthew Caesar, and P Brighten Godfrey. “Finishing flows quickly with preemptive scheduling”. In: ACM SIGCOMM Computer Communication Review 42.4 (2012), pp. 127–138.
- [HLV94] J. A. Hoogeveen, Jan Karel Lenstra, and Bart Veltman. “Three, four, five, six, or the complexity of scheduling with communication delays”. Oper. Res. Lett. 16.3 (1994), pp. 129–137.
- [HM01] C. Hanen and A. Munier. “An approximation algorithm for scheduling dependent tasks on m processors with small communication delays”. Discrete Applied Mathematics 108.3 (2001), pp. 239–257.
- [Hol02] Jonas Holmerin. “Vertex cover on 4-regular hyper-graphs is hard to approximate within $2-\epsilon$ ”. In: Proceedings on 34th Annual ACM Symposium on Theory of Computing 2002, pp. 544–552.
- [HR01] Penny E. Haxell and Vojtech Rödl. “Integer and Fractional Packings in Dense Graphs”. In: Comb. 21.1 (2001), pp. 13–38.

- [HS19] David G. Harris and Aravind Srinivasan. “The Moser-Tardos Framework with Partial Resampling”. In: *J. ACM* 66.5 (2019), 36:1–36:45.
- [HS87] Dorit S. Hochbaum and David B. Shmoys. “Using dual approximation algorithms for scheduling problems theoretical and practical results”. In: *J. ACM* 34.1 (1987), pp. 144–162.
- [HSS06] Elad Hazan, Shmuel Safra, and Oded Schwartz. “On the complexity of approximating k-set packing”. In: *Computational Complexity* 15.1 (2006), pp. 20–39.
- [Im+19] Sungjin Im, Nathaniel Kell, Janardhan Kulkarni, and Debmalya Panigrahi. “Tight Bounds for Online Vector Scheduling”. In: *SIAM J. Comput.* 48.1 (2019), pp. 93–121.
- [JCG97] Peter Jeavons, David A. Cohen, and Marc Gyssens. “Closure properties of constraints”. In: *J. ACM* 44.4 (1997), pp. 527–548.
- [Jea98] Peter Jeavons. “On the Algebraic Structure of Combinatorial Problems”. *Theor. Comput. Sci.* 200.1-2 (1998), pp. 185–204.
- [Juh82] Ferenc Juhász. “The asymptotic behaviour of Lovász’ theta-function for random graphs”. In: *Comb.* 2.2 (1982), pp. 153–155.
- [JZA19] Zhihao Jia, Matei Zaharia, and Alex Aiken. “Beyond Data and Model Parallelism for Deep Neural Networks”. In: *Proceedings of the 2nd SysML Conference, SysML ’19*. Palo Alto, CA, USA, 2019.
- [Kal04] Gil Kalai. “Social Indeterminacy”. In: *Econometrica* 72.5 (2004), pp. 1565–1581.
- [Kan+21] Dong Yeap Kang, Tom Kelly, Daniela Kuhn, Abhishek Methuku, and Deryk Osthus. “A proof of the Erdős-Faber-Lovász conjecture”. 2021. arXiv:2101.04698 [math.CO].
- [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, Geneva, Switzerland, July 9-15, 2000, Proceedings*. Vol. 1853. Lecture Notes in Computer Science. Springer, 2000, pp. 624–635.
- [Kho+07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. “Optimal Inapproximability Results for MAX-CUT and Other 2-Variable CSPs?”. In: *SIAM J. Comput.* 37.1 (2007), pp. 319–357.
- [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*, pp. 600–609.
- [Kho02a] Subhash Khot. “Hardness results for coloring 3-colorable 3-uniform hypergraphs”. In: *The 43rd Annual IEEE Symposium on Foundations of Computer Science, FOCS 2002 IEEE*. 2002, pp. 23–32.
- [Kho02b] Subhash Khot. “On the power of unique 2-prover 1-round games”. In: *Proceedings on 34th Annual ACM Symposium on Theory of Computing, STOC 2002*, 2002, pp. 767–775.

- [KMS17] Subhash Khot, Dor Minzer, and Muli Safra. “On independent sets, 2-to-2 games, and Grassmann graphs”. In: Proceedings of the 49th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2017, ACM, 2017, pp. 576–589.
- [KMS18] Subhash Khot, Dor Minzer, and Muli Safra. “Pseudorandom Sets in Grassmann Graph Have Near-Perfect Expansion”. In: 50th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, IEEE Computer Society, 2018, pp. 592–601.
- [Knu94] Donald E. Knuth. “The Sandwich Theorem”. Electron. J. Comb1 (1994).
- [KO19] Andrei A. Krokhin and Jakub Opstal. “The Complexity of 3-Colouring H-Colourable Graphs”. In: 60th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2019, IEEE Computer Society, 2019, pp. 1227–1239.
- [KP20] Jeff Kahn and Jinyoung Park. Tuza's Conjecture for random graphs. 2020. arXiv: 2007.04351 [math.CO] .
- [KR83] K. H. Kim and F. W. Roush. “On a problem of Tuza”. In: Studies in Pure Mathematics: To the Memory of Paul Turan. Basel: Birkhäuser Basel, 1983, pp. 423–425.
- [Kri95] Michael Krivelevich. “On a conjecture of Tuza about packing and covering of triangles”. In: Discret. Math.142.1-3 (1995), pp. 281–286.
- [Kro+20] Andrei A. Krokhin, Jakub Opstal, Marcin Wrochna, and Stanisław Żwiniński. “Topology and adjunction in promise constraint satisfaction”. Electron. Colloquium Comput. Complex.(2020), p. 40.
- [KS14] Subhash Khot and Rishi Saket. “Hardness of Finding Independent Sets in 2-Colorable and Almost 2-Colorable Hypergraphs”. In: Proceedings of the 25th Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2014, pp. 1607–1625.
- [KT17] Ken-ichi Kawarabayashi and Mikkel Thorup. “Coloring 3-Colorable Graphs with Less Than 1.5 Colors”. In: J. ACM64.1 (Mar. 2017).
- [Kul+20] Janardhan Kulkarni, Shi Li, Jakub Tarnawski, and Minwei Ye. “Hierarchy-Based Algorithms for Minimizing Makespan under Precedence and Communication Constraints”. In: Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, Salt Lake City, UT, USA, January 5-8, 2020, SIAM, 2020, pp. 2770–2789.
- [KZ97] Howard J. Karloff and Uri Zwick. “A $7/8$ -Approximation Algorithm for MAX 3SAT?” In: 38th Annual Symposium on Foundations of Computer Science, FOCS '97, Miami Beach, Florida, USA, October 19-22, 1997, IEEE Computer Society, 1997, pp. 406–415.
- [Law+93] Eugene L Lawler, Jan Karel Lenstra, Alexander HG Rinnooy Kan, and David B Shmoys. “Sequencing and scheduling: Algorithms and complexity”. Handbooks in operations research and management science (1993), pp. 445–522.
- [Lee19] Euiwoong Lee. “Partitioning a graph into small pieces with applications to path transversal”. In: Math. Program.177.1-2 (2019), pp. 1–19.

- [Li21] Shi Li. “Towards PTAS for Precedence Constrained Scheduling via Combinatorial Algorithms”. In: *Proceedings of the 2021 ACM-SIAM Symposium on Discrete Algorithms (SODA)*. SIAM. 2021, pp. 2991–3010.
- [LMR94] Frank Thomson Leighton, Bruce M. Maggs, and Satish Rao. “Packet Routing and Job-Shop Scheduling in $O(\text{Congestion} + \text{Dilation})$ Steps”. In: *Comb.* 14.2 (1994), pp. 167–186.
- [Lov75] László Lovász. “On minmax theorems of combinatorics, Doctoral thesis”. In: *Mathematiki Lapok* 26 (1975).
- [LR16] E. Levey and T. Rothvoss. “A $(1+\epsilon)$ -approximation for makespan scheduling with precedence constraints using LP hierarchies”. In: *Proceedings of the 48th Annual ACM SIGACT Symposium on Theory of Computing, STOC 2016, Cambridge, MA, USA, June 18-21, 2016*. 2016, pp. 168–177.
- [Luo+16] Shouxi Luo, Hongfang Yu, Yangming Zhao, Sheng Wang, Shui Yu, and Lemin Li. “Towards practical and near-optimal coflow scheduling for data center networks”. In: *IEEE Transactions on Parallel and Distributed Systems* 27.11 (2016), pp. 3366–3380.
- [LY94] Carsten Lund and Mihalis Yannakakis. “On the Hardness of Approximating Minimization Problems”. In: *J. ACM* 41.5 (1994), pp. 960–981.
- [LZ09] Linyuan Lu and Yi Zhao. “An Exact Result for Hypergraphs and Upper Bounds for the Turán Density of K_{r+1}^r ”. In: *SIAM J. Discret. Math.* 23.3 (2009), pp. 1324–1334.
- [Mai+20] Biswaroop Maiti, Rajmohan Rajaraman, David Stalfa, Zoya Svitkina, and Aravindan Vijayaraghavan. “Scheduling Precedence-Constrained Jobs on Related Machines with Communication Delay”. In: *To appear at 61st Annual IEEE Symposium on Foundations of Computer Science (FOCS 2020)* (2020).
- [Mar74] G. A. Margulis. “Probabilistic characteristics of graphs with large connectivity (in Russian)”. In: *Probl. Pered. Inform.* 10 (1974), pp. 101–108.
- [McD93] Colin McDiarmid. “A Random Recolouring Method for Graphs and Hypergraphs”. In: *Combinatorics, Probability and Computing* 2.3 (1993), pp. 363–365.
- [Mic+13] Tomasz P. Michalak, Aadithya V. Karthik, Piotr L. Szczepanski, Balaraman Ravindran, and Nicholas R. Jennings. “Efficient Computation of the Shapley Value for Game-Theoretic Network Centrality”. In: *J. Artif. Intell. Res.* 46 (2013), pp. 607–650.
- [Mir+17] Azalia Mirhoseini, Hieu Pham, Quoc V Le, Benoit Steiner, Rasmus Larsen, Yuefeng Zhou, Naveen Kumar, Mohammad Norouzi, Samy Bengio, and Jeff Dean. “Device placement optimization with reinforcement learning”. In: *Proceedings of the 34th International Conference on Machine Learning-Volume 70*. JMLR. org. 2017, pp. 2430–2439.
- [MK97] A. Munier and J.C. König. “A Heuristic for a Scheduling Problem with Communication Delays”. In: *Operations Research* 45.1 (1997), pp. 145–147.
- [MOO10] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. “Noise stability of functions with low influences: invariance and optimality”. In: *Annals of Mathematics* 171 (2010), pp. 295–341.

- [Mos10] Elchanan Mossel. “Gaussian Bounds for Noise Correlation of Functions”. In: *Geometric and Functional Analysis* 19 (2010), pp. 1713–1756.
- [Mos15] Dana Moshkovitz. “The Projection Games Conjecture and the NP-Hardness of In n-Approximating Set-Cover”. In: *Theory Comput.* 11 (2015), pp. 221–235.
- [MR10] Dana Moshkovitz and Ran Raz. “Two-query PCP with subconstant error”. In: *J. ACM* 57.5 (2010), 29:1–29:29.
- [MR95] Jiří Matoušek and Vojtěch Rödl. “On Ramsey sets in spheres”. In: *Journal of Combinatorial Theory, Series A* 70.1 (1995), pp. 30–44.
- [MRT13] Adam Meyerson, Alan Roytman, and Brian Tagiku. “Online Multidimensional Load Balancing”. In: *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2013*. 2013, pp. 287–302.
- [MS11] Monaldo Mastrolilli and Ola Svensson. “Hardness of Approximating Flow and Job Shop Scheduling Problems”. In: *J. ACM* 58.5 (2011), 20:1–20:32.
- [Nar+19] D. Narayanan, A. Harlap, A. Phanishayee, V. Seshadri, N. Devanur, G. Ganger, P. Gibbons, and M. Zaharia. “PipeDream: Generalized Pipeline Parallelism for DNN Training”. In: *Proc. 27th ACM Symposium on Operating Systems Principles (SOSP)*. Huntsville, ON, Canada, Oct. 2019.
- [Nis89] N. Nisan. “CREW PRAMs and Decision Trees”. In: *Proceedings of the Twenty-first Annual ACM Symposium on Theory of Computing*. STOC ’89. ACM, 1989, pp. 327–335.
- [NN11] Ramasuri Narayanam and Yadati Narahari. “A Shapley Value-Based Approach to Discover Influential Nodes in Social Networks”. In: *IEEE Trans Autom. Sci. Eng.* 8.1 (2011), pp. 130–147.
- [ODo14] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [OW09] Ryan O’Donnell and Yi Wu. “Conditional hardness for satisfiable 3-CSPs”. In: *Proceedings of the 41st Annual ACM Symposium on Theory of Computing*. 2009, pp. 493–502.
- [Pan+11] Rina Panigrahy, Kunal Talwar, Lincoln Uyeda, and Udi Wieder. “Heuristics for Vector Bin Packing”. 2011. URL: <https://www.microsoft.com/en-us/research/publication/heuristics-for-vector-bin-packing/>.
- [Pet20] Jan Petr. *Monotone functions avoiding majorities*. Undergraduate Thesis. Univerzita Karlova, Matematicko-fyzikalni fakulta. 2020.
- [Pip02a] Nicholas Pippenger. “Galois theory for minors of finite functions”. In: *Discrete Mathematics* 254.1-3 (2002), pp. 405–419.
- [Pip02b] Nicholas Pippenger. “Galois theory for minors of finite functions”. In: *Discrete Mathematics* 254.1 (2002), pp. 405–419.
- [Pol12] D.H.J. Polymath. “A new proof of the density Hales-Jewett theorem”. In: *Annals of Mathematics* 175.3 (May 2012), pp. 1283–1327.
- [PR81] S. Poljak and V. Rödl. “On the arc-chromatic number of a digraph”. In: *Journal of Combinatorial Theory, Series B* 31.2 (1981), pp. 190–198.

- [Pul15] Gregory J. Puleo. “Tuza’s Conjecture for graphs with maximum average degree less than 7”. In: *Eur. J. Comb.* 49 (2015), pp. 134–152.
- [PY90] Christos H. Papadimitriou and Mihalis Yannakakis. “Towards an Architecture-Independent Analysis of Parallel Algorithms”. In: *SIAM J. Comput.* 19.2 (1990), pp. 322–328.
- [Rag08] Prasad Raghavendra. “Optimal algorithms and inapproximability results for every CSP?”. In: *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, Victoria, British Columbia, Canada, May 17-20, 2008*. ACM, 2008, pp. 245–254.
- [Ray21] Arka Ray. “There is no APTAS for 2-dimensional vector bin packing: Revisited”. In: *CoRR* abs/2104.13362 (2021). arXiv: 2104. 13362.
- [Ray87] Victor J. Rayward-Smith. “UET scheduling with unit interprocessor communication delays”. In: *Discret. Appl. Math.* 18.1 (1987), pp. 55–71.
- [Raz98] Ran Raz. “A Parallel Repetition Theorem”. In: *SIAM J. Comput.* 27.3 (1998), pp. 763–803.
- [RS07] Jaikumar Radhakrishnan and Madhu Sudan. “On Dinur’s proof of the PCP theorem”. In: *Bull. Amer. Math. Soc.* 44 (2007), pp. 19–61.
- [Rus82] Lucio Russo. “An approximate zero-one law”. In: *Z. Wahrscheinlichkeitstheorie und Verwandte Gebiete* 61.1 (1982), pp. 129–139.
- [Sak14] Rishi Saket. “Hardness of Finding Independent Sets in 2-Colorable Hypergraphs and of Satisfiable CSPs”. In: *Proceedings of the 29th IEEE Conference on Computational Complexity*. 2014, pp. 78–89.
- [San21] Sai Sandeep. “Almost Optimal Inapproximability of Multidimensional Packing Problems”. In: *62nd IEEE Annual Symposium on Foundations of Computer Science, FOCS 2021, Denver, CO, USA, February 7-10, 2022*. IEEE, 2021, pp. 245–256.
- [Sch78] Thomas J. Schaefer. “The Complexity of Satisfiability Problems”. In: *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, STOC 1978*. ACM, 1978, pp. 216–226.
- [Shy+18] Ayan Shymyrbay, Arshyn Zhanbolatov, Assilkhan Amankhan, Adilya Bakambekova, and Ikechi A Ukaegbu. “Meeting Deadlines in Datacenter Networks: An Analysis on Deadline-Aware Transport Layer Protocols”. In: *2018 International Conference on Computing and Network Communications (CoCoNet)*. IEEE. 2018, pp. 152–158.
- [Sid95] Alexander Sidorenko. “What we know and what we do not know about Turán numbers”. In: *Graphs Comb.* 11.2 (1995), pp. 179–199.
- [Sid97] Alexander Sidorenko. “Upper Bounds for Turán Numbers”. In: *J. Comb. Theory, Ser. A* 77.1 (1997), pp. 134–147.
- [Sim83] Hans-Ulrich Simon. “A tight $(\log \log n)$ -bound on the time for parallel RAMs to compute nondegenerated boolean functions”. In: *Foundations of Computation Theory*. Springer Berlin Heidelberg, 1983, pp. 439–444.
- [SK67] Richard Sinkhorn and Paul Knopp. “Concerning nonnegative matrices and doubly stochastic matrices”. In: *Pacific Journal of Mathematics* 21.2 (1967), pp. 343–348.

- [Spi94] Frits CR Spieksma. “A branch-and-bound algorithm for the two-dimensional vector packing problem”. In: *Computers & operations research* 21.1 (1994), pp. 19–25.
- [SS13] Sushant Sachdeva and Rishi Saket. “Optimal Inapproximability for Scheduling Problems via Structural Hardness for Hypergraph Vertex Cover”. In: *Proceedings of the 28th Conference on Computational Complexity, CCC 2013, K.lo Alto, California, USA, 5-7 June, 2013*. 2013, pp. 219–229.
- [SS54] L. S. Shapley and Martin Shubik. “A Method for Evaluating the Distribution of Power in a Committee System”. In: *American Political Science Review* 48.3 (1954), pp. 787–792.
- [SSW94] David B. Shmoys, Clifford Stein, and Joel Wein. “Improved Approximation Algorithms for Shop Scheduling Problems”. In: *SIAM J. Comput.* 23.3 (1994), pp. 617–632.
- [ST12] Hadas Shachnai and Tami Tamir. “Approximation schemes for generalized two-dimensional vector packing with application to data placement”. In: *J. Discrete Algorithms* 10 (2012), pp. 35–48.
- [Sve10] Ola Svensson. “Conditional hardness of precedence constrained scheduling on identical machines”. In: *Proceedings of the 42nd ACM Symposium on Theory of Computing, STOC 2010, Cambridge, Massachusetts, USA, 5-8 June 2010*. ACM, 2010, pp. 745–754.
- [SW99a] P. Schuurman and G. Woeginger. “Polynomial time approximation algorithms for machine scheduling: ten open problems”. In: *Journal of Scheduling* 2.5 (1999), pp. 203–213.
- [SW99b] P. Schuurman and G. J. Woeginger. *Polynomial time approximation algorithms for machine scheduling: Ten open problems*. 1999.
- [Tar+20] Jakub Tarnawski, Amar Phanishayee, Nikhil R. Devanur, Divya Mahajan, and Fanny Nina Paravecino. *Efficient Algorithms for Device Placement of DNN Graph Operators*. 2020. arXiv: 2006. 16423 [cs. LG].
- [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*. 2001, pp. 453–461.
- [Tur41] Paul Turán. “On an extremal problem in graph theory”. In: *Mat. Fiz. Lapok* 48 (1941), pp. 436–452.
- [Tur61] Paul Turán. “Research Problem”. In: *Közl MTA Mat. Kutató Int.* 6 (1961), pp. 417–423.
- [Tuz81] Zsolt Tuza. “Conjecture”. In: *Finite and Infinite Sets Proc. Colloq. Math. Soc. Janos Bolyai* (1981), p. 888.
- [Tuz90] Zsolt Tuza. “A conjecture on triangles of graphs”. In: *Graphs Comb.* 6.4 (1990), pp. 373–380.
- [TY92] R. Thurimella and Y. Yesha. “A scheduling principle for precedence graphs with communication delay”. In: *International Conference on Parallel Processing* 3 (1992), pp. 229–236.

- [TŽ18] Johan Thapper and Stanislav Živný. “The Limits of SDP Relaxations for General-Valued CSPs”. In: *ACM Trans. Comput. Theory* 10.3 (2018), 12:1–12:22.
- [VL81] Wenceslas Fernandez de la Vega and George S. Lueker. “Bin packing can be solved within $1+\epsilon$ in linear time”. In: *Comb.* 1.4 (1981), pp. 349–355.
- [VLL90] B. Veltman, B.J. Lageweg, and J.K. Lenstra. “Multiprocessor scheduling with communication delays”. In: *Parallel Computing* 16.2 (1990), pp. 173–182.
- [Web77] Robert Weber. *Probabilistic Values for Games*. Cowles Foundation Discussion Papers 471R. Cowles Foundation for Research in Economics, Yale University, 1977.
- [Wen13] Cenny Wenner. “Circumventing d -to-1 for Approximation Resistance of Satisfiable Predicates Strictly Containing Parity of Width at Least Four”. In: *Theory of Computing* 9.23 (2013), pp. 703–757.
- [Woe97] Gerhard J. Woeginger. “There is no Asymptotic PTAS for Two-Dimensional Vector Packing”. In: *Inf. Process. Lett.* 64.6 (1997), pp. 293–297.
- [WS11] David P. Williamson and David B. Shmoys. *The Design of Approximation Algorithms*. USA: Cambridge University Press, 2011.
- [WŽ20] Marcin Wrochna and Stanislav Živný. “Improved hardness for H -colourings of G -colourable graphs”. In: *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020*. SIAM, 2020, pp. 1426–1435.
- [Yus12] Raphael Yuster. “Dense Graphs With a Large Triangle Cover Have a Large Triangle Packing”. In: *Comb. Probab. Comput.* 21.6 (2012), pp. 952–962.
- [Zha+12] Jiaxing Zhang, Hucheng Zhou, Rishan Chen, Xuepeng Fan, Zhenyu Guo, Haoxiang Lin, Jack Y Li, Wei Lin, Jingren Zhou, and Lidong Zhou. “Optimizing data shuffling in data-parallel computation by understanding user-defined functions”. In: *Presented as part of the 9th USENIX Symposium on Networked Systems Design and Implementation (NSDI 12)*. 2012, pp. 295–308.
- [Zha+15] Yangming Zhao, Kai Chen, Wei Bai, Minlan Yu, Chen Tian, Yanhui Geng, Yiming Zhang, Dan Li, and Sheng Wang. “Rapier: Integrating routing and scheduling for coflow-aware data center networks”. In: *2015 IEEE Conference on Computer Communications (INFOCOM)*. IEEE. 2015, pp. 424–432.
- [Zhu20] Dmitriy Zhuk. “A Proof of the CSP Dichotomy Conjecture”. In: *J. ACM* 67.5 (2020), 30:1–30:78.
- [Zwi98] Uri Zwick. “Finding Almost-Satisfying Assignments”. In: *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, Dallas, Texas, USA, May 23-26, 1998*. Ed. by Jeffrey Scott Vitter. ACM, 1998, pp. 551–560.