

SDPs and Robust Satisfiability of Promise CSP*

Joshua Brakensiek[†] Venkatesan Guruswami[‡] Sai Sandeep[§]

Abstract

For a constraint satisfaction problem (CSP), a robust satisfaction algorithm is one that outputs an assignment satisfying most of the constraints on instances that are near-satisfiable. It is known that the CSPs that admit efficient robust satisfaction algorithms are precisely those of bounded width, i.e., CSPs whose satisfiability can be checked by a simple local consistency algorithm (eg., 2-SAT or Horn-SAT in the Boolean case). While the exact satisfiability of a bounded width CSP can be checked by combinatorial algorithms, the robust algorithm is based on rounding a canonical Semidefinite Programming (SDP) relaxation.

In this work, we initiate the study of robust satisfaction algorithms for *promise* CSPs, which are a vast generalization of CSPs that have received much attention recently. The motivation is to extend the theory beyond CSPs, as well as to better understand the power of SDPs. We present robust SDP rounding algorithms under some general conditions, namely the existence of particular high-dimensional Boolean symmetries known as majority or alternating threshold polymorphisms. On the hardness front, we prove that the lack of such polymorphisms makes the PCSP hard for all pairs of symmetric Boolean predicates. Our approach relies on SDP integrality gaps argued via the absence of certain colorings of the sphere, with connections to sphere Ramsey theory.

We conjecture that PCSPs with robust satisfaction algorithms are precisely those for which the feasibility of the canonical SDP implies (exact) satisfiability. We also give a precise algebraic condition, known as a minion characterization, of which PCSPs have the latter property.

*An abridged preliminary version of this paper appeared in the proceedings of the 2023 Symposium on the Theory of Computing (STOC 23). This version is significantly expanded with numerous revisions and corrections.

[†]University of California, Berkeley. josh.brakensiek@berkeley.edu. Part of this work was performed while affiliated with Stanford University and supported in part by an NSF Graduate Research Fellowship and a Microsoft Research PhD Fellowship.

[‡]University of California, Berkeley. venkatg@berkeley.edu. Research supported in part by NSF grants CCF-2228287 and CCF-2211972 and a Simons Investigator award.

[§]University of California, Berkeley. saisandeep@berkeley.edu. Research supported in part by NSF grant CCF-2228287.

Contents

1	Introduction	1
1.1	Our results: Robust algorithms and hardness	3
1.2	Minion characterization of basic SDP	5
1.3	Connections to Discrete Analysis	6
2	Preliminaries	7
2.1	(P)CSP instances, templates, and polymorphisms	8
2.2	The basic SDP	12
2.3	Gram and orthogonal matrices	15
2.4	Elementary properties of Gaussians	16
3	Overview of Techniques	18
3.1	Robust algorithm for MAJ polymorphisms	18
3.2	Robust algorithm for AT polymorphisms	20
3.3	UG-hardness of robust algorithms	21
4	Robust Algorithms	23
4.1	CMM is a robust algorithm when MAJ is a polymorphism	23
4.2	Warm-up for AT: Random threshold rounding algorithm for Dual-Horn SAT	29
4.3	Algorithm for AT	31
4.4	Proof of Lemma 4.9	35
5	Unique Games based Hardness	36
5.1	Sphere Ramsey theory	39
5.2	Absence of sphere coloring via sphere Ramsey theory	42
5.3	Absence of sphere coloring via connectivity of configurations	47
6	The SDP minion	51
6.1	Minion preliminaries	51
6.2	SDP minion definition	54
6.3	On the basic SDP	55
6.4	From minion homomorphism to SDP rounding algorithm	58
6.5	From SDP rounding algorithm to minion homomorphism	59
7	Conclusion	60
A	Efficient SDP Solution	65
B	Missing Proofs	67
B.1	Missing proofs from Section 4	67
B.2	Missing proofs from Section 5	70

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: the existence of a locally consistent set of 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, 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 CSPs, namely algorithms that output solutions satisfying $1 - f(\epsilon)$ fraction of the constraints when the instance is promised to be $1 - \epsilon$ satisfiable, where $f(\epsilon) \rightarrow 0$ as $\epsilon \rightarrow 0$. Zwick obtained robust algorithms for 2-SAT using Semidefinite Programming (SDP) rounding and for Horn-SAT based on rounding a Linear Program (LP). 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. Håstad [Hås01] in his seminal work showed that even for 3-LIN (when all equations involve just three variables), for every $\epsilon, \delta > 0$, it is NP-Hard to output a solution satisfying $\frac{1}{2} + \delta$ fraction of the constraints even when the instance is guaranteed to have a solution satisfying $(1 - \epsilon)$ 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. Together with Schaefer’s dichotomy theorem [Sch78], this yields that for Boolean CSPs, bounded width characterizes robust satisfiability. For CSPs over general domains, a landmark result in the algebraic approach to CSP due to Barto and Kozik [BK14] showed that CSPs that are not bounded width can express linear equations. A reduction from Håstad’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 the class of *Promise Constraint Satisfaction Problems* (PCSPs). PCSPs are a generalization of CSPs where each constraint has a strong form and a weak form. Given an instance that is promised to have a solution satisfying the stronger form of the constraints, the objective is to find a solution satisfying the weaker form of the constraints. A classic example of PCSPs is (1-in-3-SAT, NAE-3-SAT). While both the underlying CSPs are NP-Hard, the resulting PCSP does have a polynomial-time algorithm: given an instance of 1-in-3-SAT that is promised to be satisfiable, we can find an assignment to the variables in polynomial time that satisfies each constraint as an NAE-3-SAT instance. PCSPs are a vast generalization of CSPs and capture key problems such as approximate graph and hypergraph coloring.

Since their formal introduction in [AGH17] and subsequent detailed study in [BG21] and [BBKO21], there has been a flurry of recent works on PCSPs [AB21, AD22, ABP20, BBB21, BK22, BG19, BG20, BGWŽ20, BGS23, BWŽ21, CŽ23a]. These have led to a rich and still developing theory aimed at classifying the complexity of PCSPs, by tying their (in)tractability to the symmetries associated with their defining relations, and understanding the power and limitations of various algorithmic approaches influential for CSPs in the context of PCSPs.

Against this backdrop, we initiate the study of robust algorithms for PCSPs. The motivation is two-

fold. First, as algorithms resilient to a small noise in the input, robust algorithms are important in their own right. Second, in the CSP world, the existence of efficient robust algorithms is equivalent to having bounded width and being decided by $O(1)$ levels of Sherali Adams for CSPs [TŽ17]. It is also proven [TŽ18] to be equivalent to being decided by the basic semidefinite programming (SDP) relaxation [Rag08]. Here, we say that the basic SDP decides a CSP if for every instance Φ of the CSP, Φ has an assignment satisfying all the constraints if and only if there is a vector solution satisfying all the constraints in the SDP relaxation. For CSPs, therefore, the study of robust algorithms sheds light on, and in fact, precisely captures, the power of the most popular algorithmic approaches. Robust algorithms for PCSPs provide a rich context to understand how well these algorithmic tools generalize beyond CSPs.

The main question that we are interested in this work is the following.

Question 1.1. *Which PCSPs admit polynomial time robust algorithms?*

As is the case with CSPs, a natural approach to characterize which PCSPs have robust algorithms is via the 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. Our work implies there is a robust algorithm for this PCSP. Atserias and Dalmau also proved that this PCSP is decided by $O(1)$ levels of Sherali-Adams, and as we shall prove later, it is also decided by the basic SDP.¹ On the other hand, Raghavendra’s framework of converting integrality gaps of CSPs to the hardness of approximation applies to PCSPs as well [Rag08]. In particular, his result implies that every PCSP that is not decided by the basic SDP does not admit a polynomial-time robust algorithm, assuming the Unique Games Conjecture [Kho02a]. This gives a powerful tool to show the absence of a polynomial time robust algorithm for a PCSP (albeit under the Unique Games Conjecture): showing an integrality gap for its basic SDP relaxation. With this connection, Question 1.1 naturally leads to the following question.

Question 1.2. *Which PCSPs are decided by the basic SDP relaxation?*

We make progress on Question 1.1 and Question 1.2 by studying the *polymorphisms* of PCSPs. Polymorphisms are closure properties of satisfying solutions to (Promise) CSPs. As a concrete example, consider the 2-SAT CSP: given an instance Φ of 2-SAT over n variables x_1, x_2, \dots, x_n , suppose that $\mathbf{u}, \mathbf{v}, \mathbf{w}$ are three assignments to these variables satisfying all the constraints in Φ , then the assignment \mathbf{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 in Φ . This shows that the Majority function on three variables is a polymorphism of the 2-SAT CSP. More generally, the Majority function on 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 3-LIN. On the other hand, there are no non-trivial polymorphisms for 3-SAT. Polymorphisms are the central objects in the *universal algebraic approach* to CSPs [JCG97, Jea98, BJK05, BKW17, Bul17, Zhu20], which has then been extended to PCSPs [BG21, BBKO21].

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 promise templates Γ, Γ' (that is, a set of pairs of relations) such that the polymorphism family of Γ , $\text{Pol}(\Gamma)$ is contained in $\text{Pol}(\Gamma')$, then $\text{PCSP}(\Gamma')$ (the decision problem for PCSP instances using relations from Γ') is formally easier

¹We say that the basic SDP decides a PCSP (formally defined in Section 2) if for every instance Φ of the PCSP, if there is a vector solution satisfying all the strong constraints in Φ , then, Φ has an assignment satisfying all the weak constraints.

than $\text{PCSP}(\Gamma)$, i.e., there is a *gadget reduction* from $\text{PCSP}(\Gamma')$ to $\text{PCSP}(\Gamma)$. It turns out that this gadget reduction preserves the existence of robust algorithms as well. This was used by Dalmau, Krokhin [DK13] in the CSP setting, which we extend to PCSPs in Proposition 2.10). Thus, Question 1.1 and Question 1.2 can be rephrased as *Which polymorphisms lead to robust algorithms for PCSPs? Which polymorphisms lead to being decided by the basic SDP relaxation?*

We make progress on these questions on two fronts: first, for a large class of Boolean symmetric² PCSPs where we allow negation of variables, we characterize the polymorphisms that lead to robust algorithms. Our algorithms are based on novel rounding schemes for the basic SDP relaxation, and our hardness results are proved using integrality gaps for the basic SDP relaxation. Second, towards understanding the power of basic SDP for promise CSPs, we introduce a minion \mathcal{M} and show that $\text{PCSP}(\Gamma)$ can be decided by basic SDP if and only if there is a minion homomorphism from \mathcal{M} to $\text{Pol}(\Gamma)$.

1.1 Our results: Robust algorithms and hardness

As is the case with CSPs, if a PCSP is NP-Hard, then it does not admit polynomial time robust algorithms, assuming $\text{P} \neq \text{NP}$. Thus, Question 1.1 is only relevant 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 problems based on Boolean symmetric promise templates. In [BG21], the authors showed that a Boolean symmetric PCSP with folding (i.e., we allow negating the variables) 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. Later an analogous result was shown without the folding restriction in [FKOS19] where in addition one needs to consider all families of threshold polymorphisms.

Robust algorithms. Our main algorithmic result shows that in two of these cases when the PCSP has MAJ or AT polymorphisms of all odd arities, the PCSP admits a robust algorithm. More strongly, we allow our instances to have negated variables and constant literals, which we call fiPCSP , short for *folded, idempotent PCSP*.

Theorem 1.3. *Let Γ be a Boolean promise template. If $\text{Pol}(\Gamma)$ contains AT or MAJ polymorphisms of all odd arities. Then, the robust version of $\text{fiPCSP}(\Gamma)$ admits a polynomial time algorithm. In particular,*

1. *If $\text{Pol}(\Gamma)$ contains MAJ polymorphisms of all odd arities then there exists a polynomial time algorithm that given an instance of $\text{fiPCSP}(\Gamma)$ and $\varepsilon > 0$ that is promised to have a solution satisfying $1 - \varepsilon$ fraction of the constraints, outputs a solution satisfying $1 - \tilde{O}(\varepsilon^{\frac{1}{3}})$ fraction of the constraints.³*
2. *If $\text{Pol}(\Gamma)$ contains AT polymorphisms of all odd arities then there exists a polynomial time algorithm that outputs a solution satisfying $1 - O\left(\frac{\log \log \frac{1}{\varepsilon}}{\log \frac{1}{\varepsilon}}\right)$ fraction of the constraints on an instance of $\text{fiPCSP}(\Gamma)$ promised to have a solution satisfying $1 - \varepsilon$ fraction of the constraints.*

Similar to the robust algorithms for CSPs [Zwi98, CMM09, BK16], our robust algorithms for promise templates with MAJ and AT polymorphisms are based on rounding the basic SDP relaxation. The main challenge here is to obtain robust algorithms for a large class of promise template without access to the predicates and just using the properties of their polymorphisms. We achieve this using a combination of

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, \tilde{O} hides multiplicative poly logarithmic factors.

polymorphic tools where we use the fact that the promise template contains AT or MAJ polymorphisms to deduce structural properties of the underlying predicate pairs, and SDP rounding tools where we then use these structural properties to get a robust algorithm.

For MAJ polymorphisms, we first reduce the problem to the case when every weak predicate is of the form k -SAT. We then show that the robust algorithm of Charikar, Makarychev, and Makarychev [CMM09] for the 2-SAT predicate generalizes to these classes of promise templates. While the analysis of [CMM09] is tailored towards 2-SAT, we give a completely different analysis that is not based on predicates and instead uses the existence of MAJ polymorphisms as a black box. Similarly, for the AT polymorphisms, we first use the properties of AT polymorphisms to reduce to the *weighted hyperplane* promise template that generalizes the (1-in-3-SAT, NAE-3-SAT) promise template. We then give a robust algorithm for the weighted hyperplane promise template based on a random threshold rounding technique. A detailed overview of our algorithmic ideas appears in Sections 3.1 and 3.2.

Hardness results. Unlike our robust algorithms, which work for general Boolean promise template with the said polymorphisms, in our hardness results, we rely on the symmetry of the predicates of the promise template. Furthermore, we assume that the promise template contains a single predicate pair $\Gamma = (P, Q)$ that does not admit AT or MAJ polymorphisms of all odd arities. We show that for such Boolean symmetric folded PCSPs, the basic SDP relaxation has an integrality gap with perfect completeness, i.e., there is a finite instance on which the SDP relaxation satisfies all the strong constraints with zero error but the instance is not satisfiable even using the weak constraints. By Raghavendra’s framework connecting SDP gaps and Unique-Games hardness [Rag08], the integrality gap rules out robust satisfaction algorithms (under the Unique Games conjecture (UGC) [Kho02a]).

Theorem 1.4. *Let $\Gamma = \{(P, Q)\}$ be a promise template consisting of a pair of symmetric Boolean predicates such that $\text{AT}_{L_1}, \text{MAJ}_{L_2} \notin \text{Pol}(\Gamma)$ for some odd integers L_1, L_2 . Then, under the UGC, unless $P = NP$, there is no polynomial time robust algorithm for the PCSP associated with $\text{fiPCSP}(\Gamma)$.*

Similar to our algorithmic result, we first use the properties of the polymorphisms to reduce to a small set of promise templates. We obtain integrality gaps for these fiPCSP s for the basic SDP relaxation, which then implies robust hardness under the UGC. For CSPs, strong integrality gaps [Sch08, Tul09, STT07] are known for the basic SDP relaxation and its strengthenings such as the Lasserre hierarchy, almost all of them being random constructions. For the case of PCSPs, analyzing the random constructions is trickier since we need to sample the constraints with a precise density such that there is a vector solution to the strong constraints, but the weak constraints are not satisfied. Instead, we take the opposite approach where we first construct the vector solution and then add all the constraints that the vector solution satisfies. This is similar in spirit to Feige and Schechtman’s integrality gap [FS01] for MAX-CUT where they first sampled n uniformly random points on a d -dimensional sphere and then added edges between every pair of points whose distance falls within a preset range.

In particular, we first construct an infinite integrality gap instance where the vertex set corresponds to the n -dimensional sphere \mathbb{S}^n for a large integer n and there are constraints for every tuple of vertices whose corresponding vectors satisfy the SDP constraints. For the set of promise templates that we study, we show that this instance is not satisfiable, even using weak constraints. A compactness argument then implies the existence of a finite integrality gap instance. As we shall see later, by using our minion characterization result, showing that the infinite instance has no satisfiable assignment is a necessary step to obtain a finite integrality gap instance. Toward showing that the infinite instance does not have an assignment satisfying all the weak constraints, we study colorings of the sphere $f : \mathbb{S}^n \rightarrow \{-1, +1\}$ and use a result of Matoušek and Rödl [MR95] from *sphere Ramsey theory* where the existence of monochromatic configurations in

colorings of the sphere are studied. While their result directly applies to some promise templates, for others, we combine their result with new techniques to prove the existence of structured configurations in sphere colorings. A more detailed overview appears in Section 3.3.

The power of SDPs and robust PCSP algorithms. Both our algorithmic and hardness results crucially use the basic SDP relaxation. As our algorithms for the AT and MAJ polymorphisms are based on rounding the basic SDP, we get that every Boolean promise template that contains AT or MAJ polymorphisms is decided by the basic SDP. On the hardness front, Theorem 1.4, shows that many Boolean symmetric folded PCSPs without AT or MAJ polymorphisms cannot be decided by the basic SDP. This suggests a more general relation between the basic SDP and robust algorithms for PCSPs. At an intuitive level, for both the existence of robust algorithms and being decided by the basic SDP, the underlying requirement seems to be 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 paper (noise stability is one crucial aspect that distinguishes MAJ and AT from Parity.), we believe this is a more general phenomenon and motivates us to make the following conjecture.

Conjecture 1.5. *A promise template Γ has a polynomial time robust algorithm if $\text{PCSP}(\Gamma)$ is decided by the basic SDP relaxation. Else, there is no polynomial time robust algorithm for $\text{PCSP}(\Gamma)$, unless $P = NP$.*

As mentioned earlier, if there is an integrality gap for Γ with respect to the basic SDP relaxation, then by Raghavendra’s [Rag08] result, we get that Γ does not have a polynomial time robust algorithm, assuming the Unique Games Conjecture. This already proves one direction of Conjecture 1.5. The other direction is more interesting: can we obtain robust algorithms for PCSPs just using the fact that basic SDP decides them? We remind the reader that the conjecture is already proven for CSPs, where the existence of robust algorithms [BK16] and decidability by basic SDP [TŽ18] are both shown to be equivalent to having bounded width.

1.2 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 decided via basic SDP. Originally, in the study of CSPs, such algebraic characterizations were structured as follows (e.g., [Bul17, Zhu20]).

- “Algorithm \mathcal{A} solves $\text{CSP}(\Gamma)$, if and only if there is a polymorphism $f \in \text{Pol}(\Gamma)$ with specific properties.”

Since the early days of PCSPs, it has been known that a single polymorphism cannot dictate hardness (c.f., [BG21]), and thus one must instead consider a *sequence* of polymorphisms (e.g., [BGWŽ20]):

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

However, in many cases, such a characterization is unfeasible or unwieldy. Instead, a more general approach, pioneered by [BBKO21], captures the structure of polymorphism via a *minion* (formally defined in Section 6). A key property of the polymorphisms of promise template Γ is that the function family $\text{Pol}(\Gamma)$ is closed under taking *minors*⁴. A minion is an abstraction based on this: it is a collection of objects each

⁴A function $f : D_1^n \rightarrow D_2$ of arity n is said to be a minor of another function $g : D_1^m \rightarrow D_2$ of arity m with respect to a mapping $\pi : [m] \rightarrow [n]$ such that $f(x_1, x_2, \dots, x_n) = g(x_{\pi(1)}, x_{\pi(2)}, \dots, x_{\pi(m)})$ for every $\mathbf{x} \in D_1^n$.

with an arity, and for every object a of arity m , and a mapping $\pi : [m] \rightarrow [n]$, there is a unique object b of arity n that is said to be a minor of a with respect to π . A *minion homomorphism* is a mapping between minions that preserves the minor operation. A powerful way to capture the limits of algorithms for PCSPs is via minion homomorphisms:

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

Many recent papers [BGWŽ20, CŽ23a, CŽ22] 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 1.6. *The basic SDP decides $\text{PCSP}(\Gamma)$ if and only if there is a minion homomorphism from \mathcal{M}_{SDP} to $\text{Pol}(\Gamma)$.*

We note that a similar minion was concurrently and independently discovered by Ciardo-Zivny [CŽ23b]. 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 \mathbf{v}_0 . The minors involve adding groups of vectors together. Having a minion homomorphism from \mathcal{M}_{SDP} to $\text{Pol}(\Gamma)$ implies that there are polymorphisms of Γ whose minors behave exactly like combining orthogonal vectors.

Proving Theorem 1.6 has a few technical hurdles. One challenge is that SDP solutions may require vectors of an arbitrarily large dimension. For these arbitrarily-large dimensional relationships to be captured in our minion, we have that the families of vectors making up \mathcal{M}_{SDP} reside in a (countably) infinite-dimensional vector space. Similar techniques have been used in other minion constructions [CŽ23a, CŽ22].

Another challenge that appears specifically unique to this paper is that a Basic SDP solution gives a vector corresponding to each variable, but for the proof to go through additional vectors are needed which correspond to the constraints. (The variable vectors are “projections” of the constraint vectors.) Obtaining such constraints would typically be done via Sum-of-Squares or a related routine, but we prove that including such vector constraints are without loss of generality. That is, any basic SDP solution can be extended to a solution that includes constraint vectors without modifying the original variable vectors. This gives us enough vector structure to prove that the minion homomorphism corresponds to the basic SDP solution.

Relation with sphere colorings. By a result of [BBKO21], there is a minion homomorphism $\mathcal{M}_{\text{SDP}} \rightarrow \text{Pol}(\Gamma)$ if and only if there is an assignment satisfying all the constraints in a “universal” instance of $\text{PCSP}(\Gamma)$ known as a *free structure*. In the case that Γ is a Boolean folded PCSP, this free structure for \mathcal{M}_{SDP} turns out to be an instance where every possible unit vector is a variable. The constraints correspond to collections of vectors that satisfy the corresponding basic SDP constraints. This is precisely the same infinite instance that we use to show integrality gaps. Thus, the result of [BBKO21] translates to Boolean promise templates as stating that a Boolean folded $\text{PCSP}(\Gamma)$ is decided by the basic SDP if and only if there is an assignment satisfying all the constraints in the infinite integrality gap instance. For the general theory of approximation of basic SDPs, similar constructs with sphere coloring being a ‘universal’ gap have appeared in the literature (eg. in [BHPZ21]).

1.3 Connections to Discrete Analysis

Before concluding the introduction, we want to outline some viewpoints related to Discrete Analysis, and specifically the analysis of Boolean functions, that underlie the polymorphic approach to CSPs in general, and the study of robust satisfiability in particular.

First, polymorphisms embody a discrete analogue of convexity, as they are operations under which the space of satisfying assignments to a CSP are closed. The existence of certain symmetric Boolean functions (such as OR, Majority, Parity) as polymorphisms drives the tractable cases of Boolean CSP [Sch78, Che09]. The study of Boolean promise CSPs is much richer, and includes polymorphisms such as the Alternating Threshold polymorphism which govern efficient satisfiability algorithms [BG21, BGWŻ20, BBKO21].

Among variants of promise CSPs, robust PCSPs (and more broadly approximate PCSPs) are particularly suited to the tools of discrete analysis. Satisfiability algorithms can be very brittle to noise, with Gaussian elimination to solve linear equations being the quintessential example. In fact this is inherent as solving linear systems becomes highly intractable under noise [Hås01], and this can be attributed to the fact that the underlying polymorphism (Parity function) is highly sensitive to noise. The robustness criterion seems to require that the associated polymorphisms are essentially “smooth” objects (i.e., concentrate on low-degree Fourier coefficients and demonstrate some noise stability). Understanding such phenomenon has been pivotal in developments to the theory of approximation algorithms over the last few decades, including the formulation of the Unique Games Conjecture [Kho02b, KKM07, MOO10, Mos10, O’D14, KMS18]. The more brittle algebraic methods used throughout the study of (P)CSPs therefore seem less useful. This is in part why the breakthrough of Raghavendra [Rag08] for approximate CSP satisfiability algorithms relied heavily on analytical methods whereas the proof of the CSP Dichotomy Theorem needed a number of breakthroughs in universal algebra [JCG97, Jea98, BJK05, BKW17, Bul17, Zhu20]. Even with this “smoothness” condition simplifying the landscape, proving a dichotomy for all promise CSPs will be much trickier than the corresponding result for CSPs due to a much richer family of smooth polymorphisms such as Alternating Threshold. Thus, robust algorithms for promise CSPs provides an excellent avenue to further develop the analytical machinery that has shaped the approximation algorithms field in CSPs, and potentially uncover deeper connections with universal algebra tools.

As mentioned earlier, our hardness results are obtained by connections to *sphere Ramsey theory*, and specifically the existence of certain monochromatic vector configurations in colorings of high-dimensional spheres. Such continuous analogs of questions normally posed in discrete/combinatorial settings seem interesting in their own right, and hold many exciting challenges. For instance, are there density analogs of such sphere Ramsey statements which bound the measure of subsets that avoid monochromatic configurations (see Section 7 for a discussion)?

Organization of the paper. We first start by introducing formal definitions and some general observations in Section 2 (experts familiar with SDPs and CSPs can skip or just skim this section). We then give a detailed technical overview of our results in Section 3. We provide our algorithmic results (Theorem 1.3) in Section 4 and prove the hardness results (Theorem 1.4) in Section 5. We propose and establish properties of the basic SDP minion in Section 6. Finally, we conclude in Section 7 with several intriguing challenges and open problems raised by our work.

Acknowledgment. We thank anonymous reviewers for numerous helpful comments. We thank Ryan O’Donnell and Anupam Gupta for helpful discussions on SDP solvability which led to Appendix A.

2 Preliminaries

Notations. We use $[n]$ to denote the set $\{1, 2, \dots, n\}$. We let $\mathcal{B} := \{-1, +1\}$. A predicate or a relation over a domain D of arity $k \geq 0$ is a subset of D^k . For a relation $P \subseteq D^k$ of arity k , we abuse the notation and use P both as a subset of D^k , and also as a function $P : D^k \rightarrow \{0, 1\}$. We use boldface letters to denote vectors and

roman letters to denote their elements, e.g., $\mathbf{x} = (x_1, x_2, \dots, x_k)$. We have $\mathbb{S}^n := \{\mathbf{v} \in \mathbb{R}^{n+1} : \|\mathbf{v}\|_2 = 1\}$. For a vector $\mathbf{v} \in D_1^k$ and a function $f : D_1 \rightarrow D_2$, we use $f(\mathbf{v}) \in D_2^k$ to denote $(f(v_1), f(v_2), \dots, f(v_k))$.

For a vector $\mathbf{x} \in \mathcal{B}^k$, we use $\text{hw}(\mathbf{x})$ to denote the number of $+1$ s in \mathbf{x} , i.e., $\text{hw}(\mathbf{x}) = \frac{k + \sum_{i=1}^k x_i}{2}$. For $S \subseteq \{0, 1, \dots, k\}$, we use $\text{Ham}_k S$ to denote $\{\mathbf{x} \in \{-1, +1\}^k : \text{hw}(\mathbf{x}) \in S\}$. We use NAE_k to denote the set $\text{Ham}_k\{1, 2, \dots, k-1\}$, and k -SAT to denote the set $\text{Ham}_k\{1, 2, \dots, k\}$. For vectors $\mathbf{x}, \mathbf{y} \in \mathbb{R}^n$, we use $\langle \mathbf{x}, \mathbf{y} \rangle$ to denote $\sum_i x_i y_i$.

2.1 (P)CSP instances, templates, and polymorphisms

We now work toward defining the main computational problem we seek to study: symmetric idempotent Boolean folded promise constraint satisfaction problems (PCSP) as well as the main tool we shall use for analyzing the complexity of such problems: *polymorphisms*. To build toward these definitions, we begin by defining constraint satisfaction problems (CSP). The main “parameter” defining a (P)CSP is its set of allowable constraint types, which we call a *template*. We can then use the template to build *instances* of the (P)CSP which have associated *assignments* for which one can define multiple computational *problems*.

Definition 2.1 (CSP template/instance/assignment). *We define a CSP template to be a finite set of predicates $\Gamma = \{P_1, P_2, \dots, P_l\}$ over a finite domain D , where for all $i \in [l]$, P_i has arity $k_i \geq 0$. An instance $\Phi = (V, \mathcal{C})$ of the constraint satisfaction problem (CSP) associated with the template Γ (i.e., “ Φ is an instance of $\text{CSP}(\Gamma)$ ”), consists of a variable set V and a constraint set \mathcal{C} . The variable set consists of n variables $V = \{u_1, u_2, \dots, u_n\}$. The m constraints $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ each consist of a tuple of variables $C_j = (u_{j,1}, u_{j,2}, \dots, u_{j,l_j}) \in V^{l_j}$ and an associated predicate $P^{(j)} \in \Gamma$ of the same arity l_j . An assignment $\sigma : V \rightarrow D$ is said to satisfy the constraint C_j if $\sigma(C_j) = (\sigma(u_{j,1}), \sigma(u_{j,2}), \dots, \sigma(u_{j,l_j})) \in P^{(j)}$.*

We now define two computational problems associated with $\text{CSP}(\Gamma)$ where an instance Φ of $\text{CSP}(\Gamma)$ is given as input.

Definition 2.2 (CSP decision/search problems). *In the decision version of $\text{CSP}(\Gamma)$, the objective is to decide if there is an assignment to a given instance of $\text{CSP}(\Gamma)$ that satisfies all the constraints. In the search version of $\text{CSP}(\Gamma)$, the objective is to output an assignment to a given instance of $\text{CSP}(\Gamma)$ that satisfies all the constraints.*

We next define promise constraint satisfaction problems (PCSP). We begin by defining what a promise template is.⁵

Definition 2.3 (promise template). *Let D_1, D_2 be finite domains. A promise template consists of a finite collection $\Gamma = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_l, Q_l)\}$ consists of pairs of predicates, where for all $i \in [l]$, P_i has arity $k_i \geq 0$ over domain D_1 and each Q_i has arity k_i over domain D_2 . Furthermore, we assert there exists a map $h : D_1 \rightarrow D_2$ such that for all $i \in [l]$ and $\mathbf{x} \in D_1^{k_i}$, $\mathbf{x} \in P_i$ implies $h(\mathbf{x}) \in Q_i$. In other words, h is a homomorphism from P_i to Q_i for all $i \in [l]$.*

We say that (D_1, D_2) is the domain of the promise template Γ .

Definition 2.4 (PCSP instance/assignments). *In an instance $\Phi = (V, \mathcal{C})$ of $\text{PCSP}(\Gamma)$, we have a set of n variables $V = \{u_1, u_2, \dots, u_n\}$ and m constraints $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ each consisting of a tuple of variables $C_j = (u_{j,1}, u_{j,2}, \dots, u_{j,l_j}) \in V^{l_j}$ and an associated predicate pair $(P^{(j)}, Q^{(j)}) \in \Gamma$ of the same arity l_j .*

⁵Other works, such as [BBKO21], parameterize the promise CSP as a pair of a relational structures with a homomorphism between them.

A P -assignment is a map $\sigma_1 : V \rightarrow D_1$ and a Q -assignment is a map $\sigma_2 : V \rightarrow D_2$. A P -assignment σ_1 is said to strongly satisfy the constraint C_j if $\sigma(C_j) \in P^{(j)}$, and a Q -assignment σ_2 is said to weakly satisfy the constraint C_j if $\sigma(C_j) \in Q^{(j)}$.

Crucially, due to the existence of the homomorphism h connecting each P_i with its corresponding Q_i , any P -assignment strongly satisfying a collection of constraint can be converted into a Q -assignment weakly satisfying that those constraints by composing h with the P -assignment. If $\Gamma = \{(P, Q)\}$ consists a single pair of predicates, we often state $\Gamma = (P, Q)$ and write $\text{PCSP}(P, Q)$ for $\text{PCSP}(\Gamma)$. The following are computational problems associated with instances of $\text{PCSP}(\Gamma)$.

Definition 2.5 (PCSP problems). *In the decision version of $\text{PCSP}(\Gamma)$, given an input instance $\Phi = (V, \mathcal{C})$ of $\text{PCSP}(\Gamma)$, the objective is to distinguish between the two cases.*

1. *There is a P -assignment $\sigma_1 : V \rightarrow D_1$ that strongly satisfies all the constraints.*
2. *There is no Q -assignment $\sigma_2 : V \rightarrow D_2$ that weakly satisfies all the constraints.*

In the search version of $\text{PCSP}(\Gamma)$, given an input instance $\Phi = (V, \mathcal{C})$ of $\text{PCSP}(\Gamma)$ with the promise that there is an assignment $\sigma_1 : V \rightarrow D_1$ that strongly satisfies all the constraints, the objective is to find an assignment $\sigma_2 : V \rightarrow D_2$ that weakly satisfies all the constraints.⁶

An algorithm for $\text{PCSP}(\Gamma)$ is said to be f -robust algorithm for a function $f : \mathbb{R} \rightarrow \mathbb{R}$ if the following holds. The inputs to the algorithm are an instance $\Phi = (V, \mathcal{C})$ of $\text{PCSP}(\Gamma)$ and a parameter ε with the promise that there is an assignment $\sigma_1 : V \rightarrow D_1$ that strongly satisfies $1 - \varepsilon$ fraction of the constraints. The objective of the algorithm is to output an assignment $\sigma_2 : V \rightarrow D_2$ that weakly satisfies at least $1 - f(\varepsilon)$ fraction of the constraints. We call such an algorithm a robust algorithm for $\text{PCSP}(\Gamma)$ if it is f -robust for some monotone and non-negative function f with $f(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.

In this paper, we restrict ourselves to PCSP templates where both the domains are equal to $\mathcal{B} := \{-1, +1\}$. We call such templates *Boolean* templates. For technical reasons, to more closely align with the conventions in the robust CSP literature, we consider a different notion of an PCSP instance which incorporates *variable folding*, i.e., the negation of variables. We also assume that the PCSP instance allows for the setting of constants, which we call *idempotence*.

Definition 2.6 (Boolean folded idempotent instance/assignment). *Given a Boolean template $\Gamma = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_l, Q_l)\}$ where $P_i \subseteq Q_i \subseteq \mathcal{B}^{k_i}$ for every $i \in [l]$, an instance $\Phi = (V, \mathcal{C})$ of $\text{PCSP}(\Gamma)$ consists of a set of n variables $V = \{u_1, \dots, u_n\}$ and a set of m constraints $\mathcal{C} = \{C_1, \dots, C_m\}$. For $i \in [n]$, each variable u_i has two corresponding literals: u_i and \bar{u}_i , the positive and negative literals, respectively. Observe that we identify the positive literal with the variable itself. We also have two constant literals corresponding to the elements of \mathcal{B} . For each $j \in [m]$, C_j consists of a tuple of literals $C_j = (x_{j,1}, x_{j,2}, \dots, x_{j,l_j}) \in (V \cup \bar{V} \cup \mathcal{B})^{l_j}$ and an associated predicate pair $(P^{(j)}, Q^{(j)}) \in \Gamma$ of the same arity l_j .*

Given an assignment⁷ $\sigma : V \rightarrow \mathcal{B}$, define its extended assignment $\sigma' : V \cup \bar{V} \cup \mathcal{B} \rightarrow \mathcal{B}$ to be $\sigma'(u_i) = \sigma(u_i)$ and $\sigma'(\bar{u}_i) = -\sigma(u_i)$ for every $i \in [n]$ and $\sigma'(b) = b$ for $b \in \mathcal{B}$. The assignment σ is said to strongly (and respectively weakly) satisfy the constraint C_j in Φ if $\sigma'(C_j) \in P^{(j)}$ (and respectively $\sigma'(C_j) \in Q^{(j)}$).

⁶This explains the terminology *promise* CSPs to refer to such problems. For the decision version, a promise problem consists of disjoint sets of Yes and No instances, which together need not exhaust all instances.

⁷Since the families P_i and Q_i of predicates have a common domain \mathcal{B} , we need not distinguish between “ P -assignments” and “ Q -assignments” in the Boolean setting.

We now restate the portion of Definition 2.5 for robust PCSP in the fiPCSP setting.

Definition 2.7. An algorithm for fiPCSP(Γ) is said to be f -robust algorithm for a function $f : \mathbb{R} \rightarrow \mathbb{R}$ if the following holds. The inputs to the algorithm are an instance $\Phi = (V, \mathcal{C})$ of fiPCSP(Γ) and a parameter ε with the promise that there is an assignment $\sigma_1 : V \rightarrow D_1$ that strongly satisfies $1 - \varepsilon$ fraction of the constraints. The objective of the algorithm is to output an assignment $\sigma_2 : V \rightarrow D_2$ that weakly satisfies at least $1 - f(\varepsilon)$ fraction of the constraints. We call such an algorithm a robust algorithm for fiPCSP(Γ) if it is f -robust for some monotone and non-negative function f with $f(\varepsilon) \rightarrow 0$ as $\varepsilon \rightarrow 0$.

Associated with every PCSP template, the notion of polymorphism captures the various closure properties of the satisfying solutions to any instance of the corresponding PCSP. More formally, we can define the polymorphisms of a PCSP template as follows.

Definition 2.8 (Polymorphisms of PCSPs). Given a promise template $\Gamma = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_l, Q_l)\}$ such that for every $i \in [l]$, $P_i \subseteq D_1^{k_i}$, $Q_i \subseteq D_2^{k_i}$, a polymorphism of arity L is a function $f : D_1^L \rightarrow D_2$ that satisfies the below property for every $i \in [l]$. For all $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k_i} \in D_1^{L_i}$ satisfying $((\mathbf{v}_1)_j, (\mathbf{v}_2)_j, \dots, (\mathbf{v}_{k_i})_j) \in P_i$ for each $j \in [L_i]$, we have

$$(f(\mathbf{v}_1), f(\mathbf{v}_2), \dots, f(\mathbf{v}_{k_i})) \in Q_i$$

If Γ is Boolean (i.e., $D_1 = D_2 = \mathcal{B}$), we say that $f : \mathcal{B}^n \rightarrow \mathcal{B}$ is folded if $f(-\mathbf{v}) = -f(\mathbf{v})$ for all $\mathbf{v} \in \mathcal{B}^n$. Likewise, we say that f is idempotent if $f(1, 1, \dots, 1) = 1$ and $f(-1, -1, \dots, -1) = -1$. We use $\text{Pol}(\Gamma)$ to denote the family of all the polymorphisms of Γ .

Given a predicate $P \subseteq D_1^k$ and a function $f : D_1^L \rightarrow D_2$, we let $f(P) \subseteq D_2^k$ denote the set

$$f(P) := \bigcup_{\mathbf{x}_1, \dots, \mathbf{x}_L \in P} f(\mathbf{x}_1, \dots, \mathbf{x}_L),$$

where $f(\mathbf{x}_1, \dots, \mathbf{x}_L) := (f(x_{1,1}, \dots, x_{L,1}), \dots, f(x_{1,k}, \dots, x_{L,k}))$ is the component-wise application of f . Given a family \mathcal{F} of functions $f_i : D_1^{L_i} \rightarrow D_2$ we let $O_{\mathcal{F}}(P) := \bigcup_{f \in \mathcal{F}} f(P)$.

For our algorithmic results, we focus on studying Boolean promise templates Γ whose corresponding polymorphisms contain particular families of functions. In particular, we extensively study the alternate-threshold (AT) and majority (MAJ) polymorphisms in this paper:

1. For an odd integer $L \geq 1$ and $\mathbf{x} \in \mathcal{B}^L$, we have $\text{AT}_L(\mathbf{x}) = +1$, if $x_1 - x_2 + x_3 - \dots + x_L > 0$, and -1 , otherwise.
2. For an odd integer $L \geq 1$ and $\mathbf{x} \in \mathcal{B}^L$, we have $\text{MAJ}_L(\mathbf{x}) = +1$, if $x_1 + x_2 + x_3 + \dots + x_L > 0$, and -1 , otherwise.

We let AT denote the set of AT_L for $L \geq 1$ odd and MAJ the set of MAJ_L for $L \geq 1$ odd.

Relaxations of PCSPs. We say that a promise template Γ' is a relaxation of another promise template Γ if $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma')$. If Γ' is a relaxation of Γ , then there is a *gadget reduction* from $\Gamma' \cup \{(\text{=}, \text{=})\}$ to $\Gamma \cup \{(\text{=}, \text{=})\}$, where = denotes the equality predicate in the relevant domain. More formally, it is referred to as $\Gamma' \cup \{(\text{=}, \text{=})\}$ is *positive primitive promise (ppp)-definable* from $\Gamma \cup \{(\text{=}, \text{=})\}$.

Definition 2.9 (ppp-definability of PCSPs [Che09, BG21]). We say that a pair of predicates (P', Q') of the same arity is ppp-definable from a promise template Γ over the same domain pair if there exists a fixed constant l and an instance Φ of PCSP(Γ) over $k+l$ variables $u_1, u_2, \dots, u_k, v_1, v_2, \dots, v_l$ such that

1. If $(x_1, x_2, \dots, x_k) \in P'$, then there exist y_1, y_2, \dots, y_l such that the assignment $(x_1, \dots, x_k, y_1, \dots, y_l)$ strongly satisfies all the constraints in Φ .
2. If there is an assignment (z_1, \dots, z_{k+l}) weakly satisfying all the constraints in Φ , then $(z_1, z_2, \dots, z_k) \in Q'$.

More generally, we say that a promise template Γ' is ppp-definable from another promise template Γ if every predicate pair in Γ' is ppp-definable from Γ .

Note that if Γ' is ppp-definable from Γ , then the decision version of $\text{PCSP}(\Gamma')$ can be reduced to $\text{PCSP}(\Gamma)$ in polynomial time. We now observe that the same holds for the robust version as well. More formally, we have the following proposition.⁸

Proposition 2.10. *Suppose that the promise template Γ' over (D_1, D_2) is ppp-definable from Γ over the same domain. If there exists an f -robust algorithm for $\text{PCSP}(\Gamma)$, then there exists an $O_{\Gamma, \Gamma'}(1) \cdot f$ -robust algorithm for $\text{PCSP}(\Gamma')$.*

Proof. Given an instance $\Phi' = (V', \mathcal{C}')$ of $\text{PCSP}(\Gamma')$ over a set of variables V' and containing m constraints $\mathcal{C}' = \{C'_1, C'_2, \dots, C'_m\}$, we output an instance Φ of $\text{PCSP}(\Gamma)$ containing $|V'|$ original variables and a set of dummy variables. For every constraint C'_j using (P', Q') of arity k_j involving the tuple of literals $(u_{j,1}, u_{j,2}, \dots, u_{j,k_j})$ in Φ' , we have a set of dummy variables $w_{j,1}, w_{j,2}, \dots, w_{j,l_j}$ and a set of constraints \mathcal{C}_j whose variable scope is $S_j \cup \{w_{j,1}, w_{j,2}, \dots, w_{j,l_j}\}$ as in Definition 2.9. Let $W = \cup_j \{w_{j,1}, w_{j,2}, \dots, w_{j,l_j}\}$ and let $V = V' \cup W$ be the set of variables of Φ , and $\mathcal{C} = \cup_{j \in [m]} \mathcal{C}_j$ is the set of constraints in Φ . We claim that this reduction preserves robust algorithms.

1. (Completeness). Suppose that there exists an assignment $\sigma' : V' \rightarrow D_1$ strongly satisfying $1 - \varepsilon$ fraction of the constraints in Φ' . For every $j \in [m]$ such that σ' strongly satisfies the constraint C'_j , there is an assignment $\sigma_j : \{w_{j,1}, w_{j,2}, \dots, w_{j,l_j}\} \rightarrow D_1$ such that $\sigma \cup \sigma_j : C_j \rightarrow D_1$ strongly satisfies all the constraints of \mathcal{C}_j . Consider an assignment $\sigma : V \rightarrow D_1$ where we set $\sigma(u_i) = \sigma'(u_i)$ for $u_i \in V'$, $\sigma(w_{j,i}) = \sigma_j(w_{j,i})$ for all $w_{j,i} \in W$ such that σ strongly satisfies the constraint C'_j . If σ does not strongly satisfy the constraint C'_j , we set $\sigma(w_{j,i})$ arbitrarily. The assignment σ strongly satisfies at least $1 - O_{\Gamma, \Gamma'}(\varepsilon)$ fraction of the constraints in Φ .
2. (Soundness). Suppose that there is an assignment $\sigma : V \rightarrow D_2$ weakly satisfying $1 - \varepsilon$ fraction of the constraints in Φ . By Markov's inequality, for at least $1 - O_{\Gamma, \Gamma'}(\varepsilon)$ values of $j \in [m]$, all the constraints in \mathcal{C}_j are weakly satisfied by σ . This shows that the assignment σ restricted to V' weakly satisfies $1 - O_{\Gamma, \Gamma'}(\varepsilon)$ fraction of the constraints in Φ' . \square

In the folded, idempotent setting, we also allow for negations of variables and the setting of constants, we represent these additional operations via what we coin as “fipp”-definitions.

Definition 2.11 (fipp-definability of fiPCSPs). *We say that a pair of predicates (P', Q') of the same arity is fipp-definable from a promise template Γ over the same domain pair if there exists a fixed constant l and an instance Φ of fiPCSP(Γ) over $k + l$ variables $u_1, u_2, \dots, u_k, v_1, v_2, \dots, v_l$ such that*

⁸We previously claimed falsely that the reduction holds more generally when we are given that Γ' is a relaxation of Γ . We are grateful to Andrei Krokhin for pointing this mistake to us. Under the general assumption that Γ' is a relaxation of Γ , we can reduce the robust version of Γ' to the robust version of Γ together with binary equality constraints. For the sake of exact satisfiability, the equality constraints do not matter [BKW17, BG21] but we do not know a way to obtain a robust algorithm for $\text{PCSP}(\Gamma)$ together with binary equality constraints just using the fact that Γ has a robust algorithm. Even for CSPs, a direct way to show this is not known [DK13], without using the full robust dichotomy result of [BK16].

1. If $(x_1, x_2, \dots, x_k) \in P'$, then there exist y_1, y_2, \dots, y_l such that the assignment $(x_1, \dots, x_k, y_1, \dots, y_l)$ strongly satisfies all the constraints in Φ .
2. If there is an assignment (z_1, \dots, z_{k+i}) weakly satisfying all the constraints in Φ , then $(z_1, z_2, \dots, z_k) \in Q'$.

More generally, we say that a promise template Γ' is *fipp-definable* from another promise template Γ if every predicate pair in Γ' is *fipp-definable* from Γ .

Proposition 2.12. *Suppose that the promise template Γ' over (D_1, D_2) is fipp-definable from Γ over the same domain. If there exists an f -robust algorithm for $\text{fPCSP}(\Gamma)$, then there exists an $O_{\Gamma, \Gamma'}(1) \cdot f$ -robust algorithm for $\text{fPCSP}(\Gamma')$.*

The proof of Proposition 2.12 is essentially identical to that of Proposition 2.10.

Proof. Given an instance $\Phi' = (V', \mathcal{C}')$ of $\text{fPCSP}(\Gamma')$ over a set of variables V' and containing m constraints $\mathcal{C}' = \{C'_1, C'_2, \dots, C'_m\}$, we output an instance Φ of $\text{fPCSP}(\Gamma)$ containing $|V'|$ original variables and a set of dummy variables. For every constraint C'_j using (P', Q') of arity k_j involving the tuple of literals $(u_{j,1}, u_{j,2}, \dots, u_{j,k_j})$ in Φ' , we have a set of dummy variables $w_{j,1}, w_{j,2}, \dots, w_{j,l_j}$ and a set of constraints \mathcal{C}_j whose variable scope is $S_j \cup \{w_{j,1}, w_{j,2}, \dots, w_{j,l_j}\}$ as in Definition 2.11. Let $W = \cup_j \{w_{j,1}, w_{j,2}, \dots, w_{j,l_j}\}$ and let $V = V' \cup W$ be the set of variables of Φ , and $\mathcal{C} = \cup_{j \in [m]} \mathcal{C}_j$ is the set of constraints in Φ . We claim that this reduction preserves robust algorithms.

1. (Completeness). Suppose that there exists an assignment $\sigma' : V' \rightarrow D_1$ strongly satisfying $1 - \varepsilon$ fraction of the constraints in Φ' . For every $j \in [m]$ such that σ' strongly satisfies the constraint C'_j , there is an assignment $\sigma_j : \{w_{j,1}, w_{j,2}, \dots, w_{j,l_j}\} \rightarrow D_1$ such that $\sigma \cup \sigma_j : C_j \rightarrow D_1$ strongly satisfies all the constraints of \mathcal{C}_j . Consider an assignment $\sigma : V \rightarrow D_1$ where we set $\sigma(u_i) = \sigma'(u_i)$ for $u_i \in V'$, $\sigma(w_{j,i}) = \sigma_j(w_{j,i})$ for all $w_{j,i} \in W$ such that σ strongly satisfies the constraint C'_j . If σ does not strongly satisfy the constraint C'_j , we set $\sigma(w_{j,i})$ arbitrarily. The assignment σ strongly satisfies at least $1 - O_{\Gamma, \Gamma'}(\varepsilon)$ fraction of the constraints in Φ .
2. (Soundness). Suppose that there is an assignment $\sigma : V \rightarrow D_2$ weakly satisfying $1 - \varepsilon$ fraction of the constraints in Φ . By Markov's inequality, for at least $1 - O_{\Gamma, \Gamma'}(\varepsilon)$ values of $j \in [m]$, all the constraints in \mathcal{C}_j are weakly satisfied by σ . This shows that the assignment σ restricted to V' weakly satisfies $1 - O_{\Gamma, \Gamma'}(\varepsilon)$ fraction of the constraints in Φ' . \square

2.2 The basic SDP

We now describe the Basic SDP relaxation of an instance of a PCSP, similar to how it is presented in [Rag08]. Let Γ be a promise template over domain (D, D') , and let Φ be an instance of $\text{PCSP}(\Gamma)$ over n variables $V = \{u_1, u_2, \dots, u_n\}$ and m constraints C_1, C_2, \dots, C_m . Suppose that the constraint C_j subjects the variables $(u_{j,1}, u_{j,2}, \dots, u_{j,l_j})$ to the predicate pair $(P^{(j)}, Q^{(j)})$. In the basic SDP relaxation of $\text{PCSP}(\Gamma)$ corresponding to Φ , we have a vector $\mathbf{v}_{i,a}$ corresponding to each variable u_i , $i \in [n]$, along with a label $a \in D$. We also have a unit vector \mathbf{v}_0 . For each constraint C_j , $j \in [m]$, there is a probability distribution (referred to as the *local distribution* of the constraint C_j) over the set of possible local assignments to that constraint. Since C_j may contain repeated variables, we let $S_j := \{u_{j,1}, \dots, u_{j,l_j}\}$ denote the set of variables appearing in C_j . Our probability distribution is then over the set of functions $\{f : S_j \rightarrow D\}$. We represent this using a variable $\lambda_j(f)$ for every $j \in [m]$ and local assignment $f : S_j \rightarrow D$. Note that we use $f(C_j)$ as shorthand

for $(f(u_{j,1}), f(u_{j,2}), \dots, f(u_{j,l_j}))$. Finally, we have an error parameter ε_j corresponding to the constraint j , $j \in [m]$, equal to the probability that λ_j is supported outside $P^{(j)}$. We refer to ε_j as the error of the basic SDP relaxation on the constraint j , and $\sum_{j=1}^m \varepsilon_j$ as the error of the basic SDP relaxation of the instance Φ .

$$\begin{aligned}
& \text{minimize: } \sum_{j=1}^m \varepsilon_j \\
& \text{subject to: } \varepsilon_j \geq 0 && \forall j \in [m] \\
& \lambda_j(f) \geq 0 && \forall j \in [m], f: S_j \rightarrow D \\
& \sum_{f: S_j \rightarrow D} \lambda_j(f) = 1 && \forall j \in [m] \\
& \sum_{f: S_j \rightarrow D, f(C_j) \in P^{(j)}} \lambda_j(f) = 1 - \varepsilon_j && \forall j \in [m] \\
& \langle \mathbf{v}_0, \mathbf{v}_0 \rangle = 1 \\
& \text{(First moments.) } \langle \mathbf{v}_{i,a}, \mathbf{v}_0 \rangle = \sum_{\substack{f: S_j \rightarrow D \\ f(u_i) = a}} \lambda_j(f) && \forall j \in [m], u_i \in S_j, a \in D \\
& \text{(Second moments.) } \langle \mathbf{v}_{i,a}, \mathbf{v}_{i',a'} \rangle = \sum_{\substack{f: S_j \rightarrow D \\ f(u_i) = a, f(u_{i'}) = a'}} \lambda_j(f) && \forall j \in [m], u_i, u_{i'} \in S_j, a, a' \in D
\end{aligned}$$

We implicitly assume that every variable appears in at least one clause, so the second moment condition also implies that $\langle v_{i,a}, v_{i,a'} \rangle = 0$ for $a \neq a'$ for all $i \in [n]$.

Proposition 2.13. *Let Γ be a promise template over domain (D_1, D_2) . Suppose that $\Phi = (V, \mathcal{C})$ is an instance of a PCSP(Γ) such that there is an assignment $\sigma: V \rightarrow D_1$ that strongly satisfies all the constraints. Then, the basic SDP relaxation of Φ is feasible with objective value 0.*

Proof. We embed our vectors in the real line. We set $\mathbf{v}_0 = 1$, and for every $i \in [n]$ and $a \in D_1$, we set

$$\mathbf{v}_{i,a} := \begin{cases} 1 & a = \sigma(u_i), \\ 0 & \text{otherwise.} \end{cases}$$

For any $j \in [m]$ and $f: S_j \rightarrow D_1$, we set $\lambda_j(f) = 1$ if $f(x) = \sigma(x)$ for every $x \in S_j$, and we set $\lambda_j(f) = 0$ otherwise. These variables satisfy all the constraints in the basic SDP relaxation with $\varepsilon_j = 0$ for all $j \in [m]$. \square

Boolean folded idempotent variant. In Sections 4 and 5 of the paper, we prefer to consider an alternative formulation of the Basic SDP for the Boolean folded idempotent PCSP. We let Γ be a promise template and consider an instance $\Phi = (V, \mathcal{C})$ of fiPCSP(Γ) where $V = \{u_1, u_2, \dots, u_n\}$ and $\mathcal{C} = \{C_1, C_2, \dots, C_m\}$ with the constraint C_j using the predicate pair $(P^{(j)}, Q^{(j)})$ for $j \in [m]$. Let $\mathcal{L} = V \cup \bar{V} \cup \mathcal{B}$ be the set of literals. Each clause is a tuple of literals $C_j = (x_{j,1}, x_{j,2}, \dots, x_{j,l_j}) \in \mathcal{L}^{l_j}$ bound by the predicate pair $(P^{(j)}, Q^{(j)})$.

For each variable $u_i \in V$, we have an associated vector \mathbf{v}_i along with a truth vector \mathbf{v}_0 . We extend this association to all literals by assigning a map $\mathbf{v}: \mathcal{L} \rightarrow \{\mathbf{v}_1, \dots, \mathbf{v}_n\}$ where

$$\mathbf{v}(x) = \begin{cases} \mathbf{v}_i & x = u_i \in V \\ -\mathbf{v}_i & x = \bar{u}_i \in \bar{V} \\ x \cdot \mathbf{v}_0 & x \in \mathcal{B}. \end{cases} \quad (1)$$

For each clause $C_j \in \mathcal{C}$, we let S_j denote the set of variables appearing in C_j . In particular, $S_j = \{u_i \in V : \exists k \in [l_j], x_{j,l_j} = u_i \vee x_{j,l_j} = \bar{u}_i\}$. Given a local assignment of variables $f : S_j \rightarrow \mathcal{B}$, we let $f(C_j) = (f(x_{j,1}), f(x_{j,2}), \dots, f(x_{j,l_j}))$, where we define $f(\bar{u}_i) = -f(u_i)$ and $f(\pm 1) = \pm 1$.

The relaxation is as follows.

$$\begin{aligned}
& \text{minimize: } \sum_{j=1}^m \varepsilon_j \\
& \text{subject to: } \varepsilon_j \geq 0 && \forall j \in [m] \\
& \lambda_j(f) \geq 0 && \forall j \in [m], f : S_j \rightarrow \mathcal{B} \\
& \sum_{f: S_j \rightarrow \mathcal{B}} \lambda_j(f) = 1 && \forall j \in [m] \\
& \sum_{f: S_j \rightarrow \mathcal{B}, f(C_j) \in P^{(j)}} \lambda_j(f) = 1 - \varepsilon_j && \forall j \in [m] \\
& \|\mathbf{v}_i\|_2^2 = 1 && \forall i \in \{0, 1, \dots, n\} \\
& \text{(First moments.) } \langle \mathbf{v}(x), \mathbf{v}_0 \rangle = \sum_{f: S_j \rightarrow \mathcal{B}} \lambda_j(f) f(x) && \forall j \in [m], x \in C_j \\
& \text{(Second moments.) } \langle \mathbf{v}(x), \mathbf{v}(x') \rangle = \sum_{f: S_j \rightarrow \mathcal{B}} \lambda_j(f) f(x) f(x') && \forall j \in [m], x, x' \in C_j
\end{aligned}$$

We say that basic SDP is feasible on Φ if the above objective function is zero on Φ . We show that the SDP is feasible if there is an assignment that strongly satisfies all the constraints of Φ .

Proposition 2.14. *Suppose that $\Phi = (V, \mathcal{C})$ is an instance of a Boolean folded idempotent PCSP such that there is an assignment $\sigma : V \rightarrow \mathcal{B}$ that strongly satisfies all the constraints. Then, the basic SDP relaxation of Φ is feasible with objective value 0.*

Proof. We embed our vectors in the real line. We set $\mathbf{v}_0 = 1$, and $\mathbf{v}_i := \sigma(u_i) \in \mathbb{R}$ for every $i \in [n]$. For any $j \in [m]$ and $f : S_j \rightarrow \mathcal{B}$, we set $\lambda_j(f) = 1$ if $f(x) = \sigma(x)$ for every $x \in C_j$, and we set $\lambda_j(f) = 0$ otherwise. These variables satisfy all the constraints in the basic SDP relaxation with $\varepsilon_j = 0$ for all $j \in [m]$. \square

More generally, we get that if there is an assignment that strongly satisfies $1 - \varepsilon$ fraction of the constraints in Φ , the objective value of the above relaxation is at most εm , for every $\varepsilon \geq 0$. On the other hand, if the basic SDP is feasible for an instance Φ of a PCSP, it doesn't necessarily imply that Φ has an assignment weakly satisfying all the constraints. For some PCSPs however, this is indeed the case, and we say that such PCSPs are decided by the basic SDP.

Definition 2.15. *For a promise template Γ , we say that the basic SDP decides $\text{PCSP}(\Gamma)$ if for every instance Φ such that the basic SDP is feasible on Φ , there is an assignment to Φ that weakly satisfies all the constraints.*

Efficient solvability. Polynomial-time SDP solving algorithms can only solve the problem ‘‘approximately.’’ In the SDP literature, this notion of approximation can vary widely. For our purposes, we are satisfied if the objective is approximate (computed to within $1/\text{poly}(n)$ precision), but the solution returned is a ‘‘feasible’’ solution to the SDP.

Theorem 2.16. *Suppose an instance Φ of $\text{fiPCSP}(\Gamma)$ has a solution strongly satisfying $1 - \varepsilon$ fraction of the constraints. Then, for any input $\delta > 0$, we can find a vector solution to the basic SDP whose coefficients are*

square roots of rational numbers and all inner products are rational numbers. The objective value of the SDP solution is at most $(\varepsilon + \delta)m$ and it can be computed in time $\text{poly}(n, m, 1/\delta)$.

This precise formulation does not seem to appear in the literature [GLS93, Fre04, GM12, O'D16]. As such, we show how this result can be proved using known methods. See Appendix A for details.

2.3 Gram and orthogonal matrices

Recall that a real symmetric matrix $M \in \mathbb{R}^{n \times n}$ is positive semidefinite if $\mathbf{v}^T M \mathbf{v} \geq 0$ for all $\mathbf{v} \in \mathbb{R}^n$. Since M is real and symmetric, we know that M has a full eigenspace: there exists $\lambda_1, \dots, \lambda_n \in \mathbb{R}$ and unit vectors $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^n$ such that

$$M = \sum_{i=1}^n \lambda_i \mathbf{v}_i \mathbf{v}_i^T.$$

Since M is positive semidefinite, we know that each $\lambda_i \geq 0$. In particular, if we let $V \in \mathbb{R}^{n \times n}$ be the matrix such that the i th row is $\sqrt{\lambda_i} \mathbf{v}_i$, then we have that $M = V^T V$. In other words, M is also a *Gram matrix*. It is also straightforward that any Gram matrix is positive semidefinite since if $M = V^T V$, then for any $\mathbf{v} \in \mathbb{R}^n$, we have that $\mathbf{v}^T M \mathbf{v} = \|V \mathbf{v}\|_2^2 \geq 0$. We call V a *realization* of the Gram matrix, and we call the columns of V the vectors of the realization.

A crucial fact about realizations is that they generate the same Gram matrix if and only if they are related by an *orthogonal matrix*, a matrix $Q \in \mathbb{R}^{n \times n}$ for which $Q^T Q = Q Q^T = I_n$. We prove a generalization of this equivalence where the number of vectors differs from the ambient dimension.

Proposition 2.17. *Let n, m be positive integers and let $\mathbf{u}_1, \dots, \mathbf{u}_n, \mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^m$ be vectors. The following are equivalent.*

- (1) For all $i, j \in [n]$, we have that $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \langle \mathbf{v}_i, \mathbf{v}_j \rangle$.
- (2) There exists an orthogonal matrix $Q \in \mathbb{R}^{m \times m}$ such that for all $i \in [n]$, $\mathbf{v}_i = Q \mathbf{u}_i$.

Proof. Proving that (2) implies (1) is straightforward for if $\mathbf{v}_i = Q \mathbf{u}_i$ for all $i \in [n]$, we have that

$$\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \langle Q \mathbf{u}_i, Q \mathbf{u}_j \rangle = (Q \mathbf{u}_i)^T Q \mathbf{u}_j = \mathbf{u}_i^T Q^T Q \mathbf{u}_j = \langle \mathbf{u}_i, \mathbf{u}_j \rangle,$$

where we use that Q is an orthogonal matrix.

We now prove that (1) implies (2). Let $U := \text{span}\{\mathbf{u}_i : i \in [n]\} \subset \mathbb{R}^m$. Using the Gram-Schmidt algorithm, we can construct an orthonormal basis $\mathbf{a}_1, \dots, \mathbf{a}_m$ of \mathbb{R}^m such that $\mathbf{a}_1, \dots, \mathbf{a}_{\dim U}$ is a basis of U . (Assume $\dim U \geq 1$ or else every vector equals 0 for which we can pick Q to be the identity matrix.)

Also pick $S \subseteq [n]$ such that $\{\mathbf{u}_s : s \in S\}$ is a basis of U . There exists a unique linear map $L : U \rightarrow \mathbb{R}^m$ for which $L(\mathbf{u}_s) = \mathbf{v}_s$ for all $s \in S$. We claim the following properties about L ,

- (1) $L(\mathbf{u}_i) = \mathbf{v}_i$ for all $i \in [n]$, so the image of L is $V := \text{span}\{\mathbf{v}_i : i \in [n]\}$.
- (2) $L(\mathbf{a}_1), \dots, L(\mathbf{a}_{\dim U})$ is an orthonormal basis of V .

To verify (1), for any $i \in [n]$, let $\lambda : S \rightarrow \mathbb{R}$ be such that $\mathbf{u}_i = \sum_{s \in S} \lambda(s) \mathbf{u}_s$. Then, $L(\mathbf{u}_i) = \sum_{s \in S} \lambda(s) \mathbf{v}_s$. Note then that

$$\begin{aligned} \langle L(\mathbf{u}_i), \mathbf{v}_i \rangle &= \sum_{s \in S} \lambda(s) \langle \mathbf{v}_s, \mathbf{v}_i \rangle = \sum_{s \in S} \lambda(s) \langle \mathbf{u}_s, \mathbf{u}_i \rangle = \langle \mathbf{u}_i, \mathbf{u}_i \rangle = \langle \mathbf{v}_i, \mathbf{v}_i \rangle, \text{ and} \\ \langle L(\mathbf{u}_i), L(\mathbf{u}_i) \rangle &= \sum_{s, s' \in S} \lambda(s) \lambda(s') \langle \mathbf{v}_s, \mathbf{v}_{s'} \rangle = \lambda(s) \lambda(s') \langle \mathbf{u}_s, \mathbf{u}_{s'} \rangle = \langle \mathbf{u}_i, \mathbf{u}_i \rangle = \langle \mathbf{v}_i, \mathbf{v}_i \rangle. \end{aligned}$$

This can only happen if $L(\mathbf{u}_i) = \mathbf{v}_i$. To verify (2), for each $i \in [\dim U]$, let $\kappa_i : S \rightarrow \mathbb{R}$ be such that $\mathbf{a}_i = \sum_{s \in S} \kappa_i(s) \mathbf{u}_s$. Then, note that for any $i, j \in [\dim U]$ we have that

$$\langle \mathbf{a}_i, \mathbf{a}_j \rangle = \sum_{s, s' \in S} \kappa_i(s) \kappa_j(s') \langle \mathbf{u}_s, \mathbf{u}_{s'} \rangle = \sum_{s, s' \in S} \kappa_i(s) \kappa_j(s') \langle L(\mathbf{u}_s), L(\mathbf{u}_{s'}) \rangle = \langle L(\mathbf{a}_i), L(\mathbf{a}_j) \rangle,$$

so $L(\mathbf{a}_1), \dots, L(\mathbf{a}_{\dim U})$ is an orthonormal basis of V .

To finish, we can extend $L(\mathbf{a}_1), \dots, L(\mathbf{a}_{\dim U})$ into an orthonormal basis $\mathbf{b}_1, \dots, \mathbf{b}_m$ of \mathbb{R}^m such that $\mathbf{b}_i = L(\mathbf{a}_i)$ for all $i \in [\dim U]$. We can thus construct a linear map $Q : \mathbb{R}^m \rightarrow \mathbb{R}^m$ such that $Q(\mathbf{a}_i) = \mathbf{b}_i$ for all $i \in [m]$. The matrix representation of Q is an orthogonal matrix and it sends \mathbf{u}_i to \mathbf{v}_i for all $i \in [n]$ because Q extends L . \square

As an immediate corollary, we recover the equivalence of realizations.

Corollary 2.18. *Let $M = U^T U$ and $N = V^T V$ be $n \times n$ real matrices. We have that $M = N$ if and only if there exists an orthogonal matrix Q such that $V = QU$.*

Proof. Apply Proposition 2.17 with $m = n$, $\mathbf{u}_1, \dots, \mathbf{u}_n$ as the columns of U , and $\mathbf{v}_1, \dots, \mathbf{v}_n$ as the columns of V . \square

2.4 Elementary properties of Gaussians

We state 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 2.19. *Suppose that $X \sim \mathcal{N}(0, \sigma^2)$ is a Gaussian random variable with zero mean and variance σ^2 . Then, for every $c \in \mathbb{R}$ and $\varepsilon \geq 0$,*

$$\Pr(X \in [c - \varepsilon, c + \varepsilon]) \leq \frac{\varepsilon}{\sigma}.$$

Proof. By the definition of the Gaussian density, we have:

$$\Pr(X \in [c - \varepsilon, c + \varepsilon]) = \int_{c - \varepsilon}^{c + \varepsilon} \frac{1}{\sqrt{2\pi\sigma^2}} e^{-\frac{(x-\mu)^2}{2\sigma^2}} dx.$$

By substituting $Z = \frac{X}{\sigma}$, we get

$$\begin{aligned} \Pr(X \in [c - \varepsilon, c + \varepsilon]) &= \Pr\left(Z \in \left[\frac{c - \varepsilon}{\sigma}, \frac{c + \varepsilon}{\sigma}\right]\right) \\ &= \int_{\frac{c - \varepsilon}{\sigma}}^{\frac{c + \varepsilon}{\sigma}} \frac{1}{\sqrt{2\pi}} e^{-\frac{z^2}{2}} dz \\ &\leq \frac{2\varepsilon}{\sqrt{2\pi}\sigma} \\ &\leq \frac{\varepsilon}{\sigma}. \end{aligned} \quad \square$$

Proposition 2.20. *Suppose that $X \sim \mathcal{N}(0, 1)$ has the standard Gaussian distribution. Then, for every c and $\varepsilon \geq 0$,*

$$\Pr(X \in [c - \varepsilon, c + \varepsilon]) \leq \varepsilon.$$

Proof. We have

$$\Pr(X \in [c - \varepsilon, c + \varepsilon]) = \int_{c-\varepsilon}^{c+\varepsilon} \frac{1}{\sqrt{2\pi}} e^{-\frac{x^2}{2}} dx \leq \int_{c-\varepsilon}^{c+\varepsilon} \frac{1}{\sqrt{2\pi}} dx \leq \varepsilon. \quad \square$$

We also need the following concentration inequality for 1-dimensional Gaussian.

Proposition 2.21. [Ver18] Suppose that $X \sim \mathcal{N}(0, \sigma^2)$ has Gaussian distribution with variance σ^2 . Then, for every $t \geq 0$,

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

We will need the following properties of the multivariate Gaussian distribution:

Proposition 2.22. Let $\zeta \sim \mathcal{N}(\mathbf{0}, \Sigma)$ be an n -dimensional multivariate Gaussian random vector.

1. For any fixed vector $\mathbf{v} \in \mathbb{R}^n$, the random variable $\langle \zeta, \mathbf{v} \rangle$ follows $\mathcal{N}(0, \mathbf{v}^T \Sigma \mathbf{v})$.
2. For any two fixed vectors $\mathbf{v}, \mathbf{w} \in \mathbb{R}^n$, the covariance of the random variables $\langle \zeta, \mathbf{v} \rangle$ and $\langle \zeta, \mathbf{w} \rangle$ is equal to $\mathbf{v}^T \Sigma \mathbf{w}$.

Proof. The random variable $r = \langle \zeta, \mathbf{v} \rangle = \sum_{i \in [n]} v_i \zeta_i$ is a linear combination of Gaussian random variables, and therefore, follows a Gaussian distribution. We calculate its mean and variance as follows:

$$\mathbb{E}[r] = \mathbb{E} \left[\sum_{i \in [n]} v_i \zeta_i \right] = \sum_{i \in [n]} v_i \mathbb{E}[\zeta_i] = 0.$$

The variance of r is:

$$\text{Var}(r) = \mathbb{E}[r^2] = \mathbb{E} \left[\left(\sum_{i \in [n]} v_i \zeta_i \right)^2 \right] = \sum_{i, j \in [n]} v_i v_j \mathbb{E}[\zeta_i \zeta_j].$$

Since $\mathbb{E}[\langle \zeta, \zeta^T \rangle] = \Sigma$,

$$\text{Var}(r) = \sum_{i, j \in [n]} v_i v_j \Sigma_{ij} = \mathbf{v}^T \Sigma \mathbf{v}.$$

Therefore, $r \sim \mathcal{N}(0, \mathbf{v}^T \Sigma \mathbf{v})$, which proves the first part.

Let $r' = \langle \zeta, \mathbf{w} \rangle = \sum_{i \in [n]} w_i \zeta_i$. The covariance of r and r' is:

$$\text{Cov}(r, r') = \mathbb{E}[r r'] = \mathbb{E} \left[\left(\sum_{i \in [n]} v_i \zeta_i \right) \cdot \left(\sum_{j \in [n]} w_j \zeta_j \right) \right].$$

Expanding the product, we have:

$$\mathbb{E}[r r'] = \sum_{i, j \in [n]} v_i w_j \mathbb{E}[\zeta_i \zeta_j] = \sum_{i, j \in [n]} v_i w_j \Sigma_{ij}.$$

This simplifies to:

$$\text{Cov}(r, r') = \mathbf{v}^T \Sigma \mathbf{w}. \quad \square$$

3 Overview of Techniques

Moving forward, unless otherwise specified, all promise templates Γ are Boolean, and we consider the folded, idempotent variant of promise constraint satisfaction, i.e., $\text{fiPCSP}(\Gamma)$.

3.1 Robust algorithm for MAJ polymorphisms

We obtain our robust algorithms by first solving the basic SDP relaxation and then *rounding* the vectors. We first illustrate the SDP rounding idea with a warm-up algorithm (originally appeared in [Zwi98]) to solve the decision version of 2-SAT CSP predicate $P = \{(-1, +1), (+1, -1), (+1, +1)\}$. Consider an instance Φ of 2-SAT (i.e., $\text{fiPCSP}(P, P)$) over a set of n variables and m constraints. We first solve the basic SDP relaxation of Φ . If the basic SDP relaxation has a strictly positive error, then the instance is clearly not satisfiable. For the moment, suppose that the relaxation has zero error and we found a set of vectors $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_n$, and the local probability variables $\lambda_j(f)$ for all $j \in [n], f : S_j \rightarrow \mathcal{B}$ that satisfy all the constraints in the basic SDP relaxation with $\varepsilon_j = 0$ for every $j \in [m]$. Consider an arbitrary constraint C_j using the literals x_1, x_2 . We abuse the notation and let $\mathbf{v}_1, \mathbf{v}_2$ denote the vectors assigned by the SDP to x_1 and x_2 respectively (i.e., $\mathbf{v}(x_1)$ and $\mathbf{v}(x_2)$). Using the first moment and second moment properties satisfied by the local probabilities $\lambda_j(f)$, we get the following.

1. $\langle \mathbf{v}_1, \mathbf{v}_0 \rangle + \langle \mathbf{v}_2, \mathbf{v}_0 \rangle \geq 0$.
2. If $\langle \mathbf{v}_1, \mathbf{v}_0 \rangle + \langle \mathbf{v}_2, \mathbf{v}_0 \rangle = 0$, $\mathbf{v}_1 + \mathbf{v}_2 = 0$.

These properties motivate the following simple rounding algorithm that outputs the assignment $\sigma : V \rightarrow \mathcal{B}$. We sample a random vector $\zeta \sim \mathcal{N}(0, \mathbf{I})$, and set

$$\sigma(u_i) = \begin{cases} +1, & \text{if } \langle \mathbf{v}_i, \mathbf{v}_0 \rangle > 0. \\ -1, & \text{if } \langle \mathbf{v}_i, \mathbf{v}_0 \rangle < 0. \\ +1, & \text{if } \langle \mathbf{v}_i, \mathbf{v}_0 \rangle = 0 \text{ and } \langle \zeta, \mathbf{v}_i \rangle > 0. \\ -1, & \text{if } \langle \mathbf{v}_i, \mathbf{v}_0 \rangle = 0 \text{ and } \langle \zeta, \mathbf{v}_i \rangle < 0. \end{cases}$$

This ensures that if a constraint uses the literals x_1, x_2 , at least one of x_1 or x_2 is rounded to $+1$.

While this algorithm finds a satisfying solution when the underlying instance Φ has a solution satisfying all the constraints, it does not give any non-trivial guarantees when the instance is only promised to have a solution satisfying $1 - \varepsilon$ fraction of the constraints. By setting $\delta = \varepsilon$ in Theorem 2.16, we get that the average error of the basic SDP relaxation is at most $O(\varepsilon)$ on any instance that is promised to satisfy $1 - \varepsilon$ fraction of the constraints. Zwick [Zwi98] gave a robust algorithm based on rounding the basic SDP relaxation for 2-SAT which was later improved by Charikar, Makarychev, and Makarychev [CMM09]. They sample a random vector $\zeta \sim \mathcal{N}(0, \mathbf{I})$, and set

$$\sigma(u_i) = \begin{cases} +1, & \text{if } \langle \mathbf{v}_i, \zeta \rangle \geq -\frac{\langle \mathbf{v}_i, \mathbf{v}_0 \rangle}{\sqrt{\varepsilon}}. \\ -1, & \text{otherwise.} \end{cases}$$

One can view their algorithm as a smoothed version of the earlier discussed algorithm: Using Proposition 2.22, the random variable $r_i = \langle \mathbf{v}_i, \zeta \rangle$ follows $\mathcal{N}(0, 1)$ for all $i \in [n]$. Thus, if $|\langle \mathbf{v}_i, \mathbf{v}_0 \rangle| > \lambda\sqrt{\varepsilon}$ for a positive real number λ , then $\sigma(u_i)$ is equal to the sign of $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle$ if $r_i \leq \lambda$, which occurs with probability

at least $1 - e^{-\frac{\lambda^2}{2}}$ (Proposition 2.21). On the other hand, if $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = 0$, we set $\sigma(u_i) = +1$ if and only if $\langle \zeta, \mathbf{v}_i \rangle \geq 0$.

We use this algorithm in our proof of Theorem 1.3 but we give a completely different analysis. Consider a constraint C_j in the 2-SAT instance and let $\mathbf{v}_1, \mathbf{v}_2$ denote the vectors assigned by the basic SDP to the literals in a constraint. Let $r_1 = \langle \mathbf{v}_1, \zeta \rangle$ and $r_2 = \langle \mathbf{v}_2, \zeta \rangle$. From Proposition 2.22, both the random variables r_1 and r_2 are standard Gaussian variables with covariance $\text{Cov}(r_1, r_2) = \langle \mathbf{v}_1, \mathbf{v}_2 \rangle$. To upper bound the probability that the output assignment σ violates the constraint C_j , [CMM09] calculate the probability that $r_1 < -\frac{\langle \mathbf{v}_1, \mathbf{v}_0 \rangle}{\sqrt{\varepsilon}}$ and $r_2 < -\frac{\langle \mathbf{v}_2, \mathbf{v}_0 \rangle}{\sqrt{\varepsilon}}$. They do so by computing the probability that a pair of Gaussian random variables with known covariance lie in a given intersection of two half-spaces. While this analysis works for 2-SAT, these calculations turn out to be significantly harder when there are more than two Gaussian random variables.

Instead, we take a conceptually different, and arguably simpler route. As a concrete example, consider the promise template $\Gamma = (P, Q)$ where $P = \text{Ham}_4\{2, 3, 4\}$, $Q = \text{Ham}_4\{1, 2, 3, 4\}$, i.e.,

$$P = \{\mathbf{x} \in \mathcal{B}^4 : \sum_{i \in 4} x_i \geq 0\}, \quad Q = \mathcal{B}^k \setminus \{(-1, -1, -1, -1)\}.$$

Consider a constraint C_j of an instance Φ of $\text{fPCSP}(\Gamma)$. As the average error of SDP over all the constraints is at most ε , for at least $1 - \varepsilon^{1/4}$ fraction of the constraints, the error is at most $\varepsilon^{3/4}$. We restrict ourselves to these constraints and let $\mathbf{v}_1, \dots, \mathbf{v}_4$ denote the vectors assigned by the SDP to the literals in the constraint. We have that $\sum_{i \in [4]} \langle \mathbf{v}_i, \mathbf{v}_0 \rangle \geq -\varepsilon^{3/4}$. Recall that if for some $i \in [4]$, $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle \geq C\sqrt{\varepsilon}$ for a positive real number C , the corresponding literal is rounded to $+1$ with probability at least $1 - e^{-\frac{C^2}{2}}$. On the other hand, if for some $i \in [4]$, $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle \leq -4C\sqrt{\varepsilon}$, there is some $i' \in [4]$ with $\langle \mathbf{v}_{i'}, \mathbf{v}_0 \rangle \geq C\sqrt{\varepsilon}$, which again ensures that there is at least one literal that is rounded to $+1$ with probability at least $1 - e^{-\frac{C^2}{2}}$. Hence, if there is at least one $i \in [4]$ with $|\langle \mathbf{v}_i, \mathbf{v}_0 \rangle| \geq 4C\sqrt{\varepsilon}$, at least one literal is rounded to $+1$ with probability at least $1 - e^{-\frac{C^2}{2}}$.

Thus the interesting case is when $|\langle \mathbf{v}_i, \mathbf{v}_0 \rangle| \leq O(\sqrt{\varepsilon})$ for every $i \in [4]$. In this case, using the first and second moment properties satisfied by these vectors, we get that $\|\sum_{i \in [4]} \mathbf{v}_i\|_2$ is at most $O(\varepsilon^{1/4})$. The output assignment σ violates Q on this constraint only if $\langle \mathbf{v}_i, \zeta \rangle \leq -\frac{\langle \mathbf{v}_i, \mathbf{v}_0 \rangle}{\sqrt{\varepsilon}}$ for every $i \in [4]$, or equivalently, $\langle \mathbf{v}_i, \zeta \rangle + \frac{\langle \mathbf{v}_i, \mathbf{v}_0 \rangle}{\sqrt{\varepsilon}} \leq 0$ for every $i \in [4]$. However, using $\|\sum_{i \in [4]} \mathbf{v}_i\|_2 \leq O(\varepsilon^{1/4})$ and $\sum_{i \in [4]} \langle \mathbf{v}_i, \mathbf{v}_0 \rangle \geq -\varepsilon^{3/4}$, we get that

$$\sum_{i \in [4]} \left(\langle \mathbf{v}_i, \zeta \rangle + \frac{\langle \mathbf{v}_i, \mathbf{v}_0 \rangle}{\sqrt{\varepsilon}} \right) \geq -O(\varepsilon^{1/4}).$$

If σ does not satisfy Q , then for some $i \in [4]$, we have that

$$\langle \mathbf{v}_i, \zeta \rangle + \frac{\langle \mathbf{v}_i, \mathbf{v}_0 \rangle}{\sqrt{\varepsilon}} \in \left[-O(\varepsilon^{1/4}), 0 \right]$$

Finally, we can upper bound the probability that this occurs to be at most $\tilde{O}(\varepsilon^{1/4})$ using anti-concentration of the Gaussian $\langle \mathbf{v}_i, \zeta \rangle \sim \mathcal{N}(0, 1)$ (Proposition 2.20). Thus, we can infer that the assignment σ satisfies at least $1 - \tilde{O}(\varepsilon^{1/4})$ fraction of the constraints in expectation. A careful analysis of the parameters gives a guarantee of $1 - \tilde{O}(\varepsilon^{1/3})$.

For an arbitrary Boolean promise template Γ with $\text{MAJ} \subseteq \text{Pol}(\Gamma)$, we obtain a robust algorithm for $\text{fPCSP}(\Gamma)$ by first reducing to the case when all the predicate pairs are of the form (P, Q) with $Q = \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\}$, generalizing the above two examples of the 2-SAT template (i.e., $P = Q = \{(-1, +1), (1, -1), (1, 1)\}$ and $(\text{Ham}_4\{2, 3, 4\}, \text{Ham}_4\{1, 2, 3, 4\})$). Then, we find a weight vector \mathbf{w} which satisfies that $\langle \mathbf{w}, \mathbf{x} \rangle \geq 0$ for all $\mathbf{x} \in P$ ($\mathbf{w} = (1, 1, \dots, 1)$ suffices for the previous two examples). We prove

the existence of such a vector \mathbf{w} by using a Linear Programming relaxation, and we crucially use the fact that $\text{MAJ} \subseteq \text{Pol}(P, Q)$ in the analysis of this LP relaxation. Once we find the vector \mathbf{w} , the above analysis of $(\text{Ham}_4\{2, 3, 4\}, \text{Ham}_4\{1, 2, 3, 4\})$ can be generalized, the main change being that we study the properties of the weighted sum of the \mathbf{v}_i s with weights being given by the vector \mathbf{w} .

3.2 Robust algorithm for AT polymorphisms

For the Alternating-Threshold (AT) case, we combine these ideas with a random geometric sampling trick. As a concrete example, consider the templates for 1-in-3-SAT and NAE-SAT: $P = \{\mathbf{x} \in \mathcal{B}^3 : x_1 + x_2 + x_3 = -1\}$ and $Q = \{\mathbf{x} \in \mathcal{B}^3 : x_1 + x_2 + x_3 \in \{-1, 1\}\}$, respectively. We have that $\text{AT} \subseteq \text{Pol}(P, Q)$ [BG21]. Assume that an instance (V, \mathcal{C}) of $\text{fPCSP}(P, Q)$ is perfectly satisfiable. For this case, we can efficiently find an exact solution using the basic SDP relaxation via random hyperplane rounding as follows. Consider an arbitrary constraint C_j and let $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ denote the vectors assigned by the basic SDP to the literals in C_j . Using the fact that these vectors satisfy the first and second moment constraints of the basic SDP relaxation with zero error, we can infer that their sum $\mathbf{v}_s = \mathbf{v}_1 + \mathbf{v}_2 + \mathbf{v}_3$ is equal to $-\mathbf{v}_0$ for every constraint C . Let $\mathbf{v}_i \perp \mathbf{v}_0 = \mathbf{v}_i - \langle \mathbf{v}_i, \mathbf{v}_0 \rangle \mathbf{v}_0$ for $i \in [3]$. Note that $\sum_{i \in [3]} \mathbf{v}_i \perp \mathbf{v}_0 = 0$. Using this observation, we can design a rounding scheme. We first sample $\zeta \sim \mathcal{N}(0, \mathbf{I})$, and set $\sigma(u_i) = +1$ if $\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle > 0$, and -1 if $\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle < 0$. Each variable u_i is assigned a value with probability 1 unless $\mathbf{v}_i \in \{\mathbf{v}_0, -\mathbf{v}_0\}$ since \mathbf{v}_i is a unit vector. In that undetermined case, we set $\sigma(\mathbf{u}_i) = +1$ if $\mathbf{v}_i = \mathbf{v}_0$ and $\sigma(\mathbf{u}_i) = -1$ otherwise.⁹ As $\mathbf{v}_1 + \mathbf{v}_2 + \mathbf{v}_3 = -\mathbf{v}_0$, the rounding scheme ensures that at least one literal associated with these vectors is set to $+1$, and at least one literal is set to -1 .

For the robust setting where we are only guaranteed that there is an assignment strongly satisfying $1 - \varepsilon$ fraction of the constraints. By setting $\delta = \varepsilon$ in Theorem 2.16, we get that the average SDP error is at most $O(\varepsilon)$. By Markov's inequality, we are guaranteed that for at least $1 - \sqrt{\varepsilon}$ fraction of the constraints, the SDP error is at most $O(\sqrt{\varepsilon})$. For these constraints, we get that the sum vector \mathbf{v}_s 's component orthogonal to \mathbf{v}_0 has ℓ_2 norm at most $O(\varepsilon^{1/4})$, i.e., $\|\sum_{i \in [3]} \mathbf{v}_i \perp \mathbf{v}_0\|_2 \leq O(\varepsilon^{1/4})$. Using this, we design a rounding scheme that is similar to the above, with the addition that when $\|\mathbf{v}_i \perp \mathbf{v}_0\|_2$ is very small, we want to round it to $+1$ or -1 depending on its component along \mathbf{v}_0 , similar to how we were rounding \mathbf{v}_0 to $+1$ and $-\mathbf{v}_0$ to -1 in the exact algorithm earlier. We have the following compact algorithm based on this idea.

$$\sigma(u_i) = \begin{cases} +1, & \text{if } \langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle \geq -\delta \langle \mathbf{v}_i, \mathbf{v}_0 \rangle. \\ -1, & \text{otherwise.} \end{cases}$$

We choose δ to be ε^κ , where κ is a constant such that $\kappa < \frac{1}{4}$, to be set later. Note that when $|\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle| > \delta$, our new rounding scheme is equivalent to the earlier exact algorithm, i.e., $\sigma(u_i) = +1$ if $\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle > 0$, and -1 otherwise.

If $|\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle| > 3\delta$ for some $i \in [3]$, the condition $\|\sum_{i \in [3]} \mathbf{v}_i \perp \mathbf{v}_0\|_2 \leq O(\varepsilon^{1/4}) < \delta$ implies that there exist distinct $i, i' \in [3]$ such that $\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle > \delta$ and $\langle \mathbf{v}_{i'} \perp \mathbf{v}_0, \zeta \rangle < -\delta$. In this case, $\sigma(u_i) = +1$ and $\sigma(u_{i'}) = -1$, ensuring that the output assignment satisfies the NAE-3-SAT constraint. On the other hand, when $|\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle|$ is much smaller than δ for every $i \in [3]$, our rounding function sets $\sigma(u_i) = +1$ if $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle > 0$ and -1 otherwise. Since $\sum_{i \in [3]} \mathbf{v}_i$ is close to $-\mathbf{v}_0$, the output satisfies the NAE-3-SAT constraint even in this case.

Finally, we employ a geometric sampling trick where δ is sampled uniformly at random from a geometric series (by choosing κ from an arithmetic series) to ensure that, with high probability, either $|\langle \mathbf{v}_i \perp$

⁹Observe this choice of assignment is consistent with the default rounding of the constant literals ± 1 .

$\mathbf{v}_0, \zeta) > 3\delta$ for some $i \in [3]$, or $|\langle \mathbf{v}_i \perp \mathbf{v}_0, \zeta \rangle|$ is much smaller than δ for every $i \in [3]$.

For an arbitrary Boolean promise template Γ with $\text{AT} \subseteq \text{Pol}(\Gamma)$, we first show that $\text{fiPCSP}(\Gamma)$ reduces to $\text{fiPCSP}(P, Q)$ where

$$P = \{\mathbf{v} \in \mathcal{B}^k : \langle \mathbf{v}, \mathbf{w} \rangle = b\}, \quad Q = \mathcal{B}^k \setminus \{\mathbf{x}, -\mathbf{x}\}$$

for some $\mathbf{w} \in \mathbb{R}^k, b \in \mathbb{R}$ with $w_i \neq 0 \forall i \in [k]$, where we set \mathbf{x} as $x_i = +1$ if $w_i > 0$, and -1 otherwise. This generalizes our previous 1-in-3-SAT versus NAE-3-SAT template which corresponds to the case when $\mathbf{w} = (1, 1, 1), b = -1$. We show that the algorithm above works for this general predicate pair as well, thereby obtaining for every promise template Γ with $\text{AT} \subseteq \text{Pol}(\Gamma)$ a robust algorithm for $\text{fiPCSP}(\Gamma)$.

3.3 UG-hardness of robust algorithms

As mentioned earlier, our Unique Games hardness for $\text{fiPCSP}(\Gamma)$ (Theorem 1.4) is based on an integrality gap for the basic SDP relaxation, i.e., we need to show that there is a finite instance Φ of $\text{fiPCSP}(\Gamma)$ that has SDP error of zero, yet there is no assignment that weakly satisfies all the constraints. In pursuit of this goal, we develop a general recipe for showing integrality gaps with respect to basic SDP for Promise CSPs via colorings of the n -dimensional unit sphere \mathbb{S}^n .

We focus on Boolean promise templates which are symmetric. We say that a predicate $P \subseteq \mathcal{B}^k$ is symmetric if for every $\mathbf{x}, \mathbf{y} \in \mathcal{B}^k$ such that $\text{hw}(\mathbf{x}) = \text{hw}(\mathbf{y})$, we have $\mathbf{x} \in P$ if and only if $\mathbf{y} \in P$.

We first start by showing an integrality gap instance for 3-LIN. Recall that the 3-LIN CSP has the predicate P with

$$P = \{\mathbf{x} \in \mathcal{B}^3 : x_1 + x_2 + x_3 = -1 \text{ or } x_1 + x_2 + x_3 = +3\}. \quad (2)$$

Consider the instance Φ of $\text{fiPCSP}(P, P)$ that uses three variables and uses two constraints C_1, C_2 with $C_1 = \{(x_1, x_2, x_3)\}, C_2 = \{(\bar{x}_1, \bar{x}_2, \bar{x}_3)\}$. The instance has no assignment that satisfies both the constraints, and we now show that the basic SDP solution has zero error on Φ . We first describe the local probability variables: for each $j \in [2]$ and each assignment $f : S_j \rightarrow \mathcal{B}$, where $S_j = \{u_1, u_2, u_3\}$, we set

$$\lambda_j(f) = \begin{cases} 1/4 & f(C_j) \in P \\ 0 & \text{otherwise.} \end{cases}$$

That is, we set each local distribution to be the uniform distribution over satisfying assignments. Substituting these in the first and second moment constraints, we get the following requirements that the vectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ need to satisfy.

$$\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = 0 \quad \forall i \in [3] \quad (3)$$

$$\langle \mathbf{v}_i, \mathbf{v}_{i'} \rangle = 0 \quad \forall i \neq i' \in [3] \quad (4)$$

We can find such three vectors by picking three orthogonal vectors that are all orthogonal to \mathbf{v}_0 . This shows that there is a solution to the basic SDP relaxation of Φ that has zero error, thus finishing the proof of the existence of an integrality gap for Φ .

While the simple example gives an integrality gap for 3-LIN, it is a challenging task to find such explicit integrality gap instances for general predicates. We develop a non-explicit approach where for a given promise template Γ we first construct an infinite integrality gap instance $\mathcal{I}^n(\Gamma)$ for $\text{fiPCSP}(\Gamma)$ and then use it to show the existence of a finite integrality gap instance. The variable set V of $\mathcal{I}^n(\Gamma)$ to be the set of unit

vectors in \mathbb{R}^{n+1} : $V = \{u_{\mathbf{v}} : \mathbf{v} \in \mathbb{S}^n\}$. We fix an arbitrary vector to be assigned \mathbf{v}_0 . We let $\mathcal{L} = V \cup \bar{V} \cup \mathcal{B}$ be our set of literals, and let $\hat{\mathbf{v}} : \mathcal{L} \rightarrow \mathbb{S}^n$ be the map

$$\hat{\mathbf{v}}(x) = \begin{cases} \mathbf{v} & x = u_{\mathbf{v}} \\ -\mathbf{v} & x = \bar{u}_{\mathbf{v}} \\ x \cdot \mathbf{v}_0 & x \in \mathcal{B}. \end{cases}$$

We now describe the clauses \mathcal{C} of our integrality gap instance. Consider $(P, Q) \in \Gamma$ of arity k . For a potential clause $C = (x_1, \dots, x_k) \in \mathcal{L}^k$, let S be the set of variables making up C . We add C to \mathcal{C} if there exists a probability distribution λ over assignments $f : S(C) \rightarrow \mathcal{B}$ satisfying the following constraints.

$$\begin{aligned} & \sum_{f: S \rightarrow \mathcal{B}, f(C) \in P} \lambda_j(f) = 1 \\ \text{(First moments.)} \quad & \langle \hat{\mathbf{v}}(x), \mathbf{v}_0 \rangle = \sum_{f: S \rightarrow \mathcal{B}} \lambda(f) f(x) \quad x \in C \\ \text{(Second moments.)} \quad & \langle \hat{\mathbf{v}}(x), \hat{\mathbf{v}}(x') \rangle = \sum_{f: S \rightarrow \mathcal{B}} \lambda(f) f(x) f(x') \quad x, x' \in C. \end{aligned}$$

We refer to $\hat{\mathbf{v}}(C) := (\hat{\mathbf{v}}(x_1), \dots, \hat{\mathbf{v}}(x_k))$ satisfying the above as a P -configuration with respect to \mathbf{v}_0 . The way we have added the constraints ensures that setting \mathbf{a} to the variable $u_{\mathbf{a}}$ satisfies all the basic SDP constraints with zero error. We then show that the instance $\mathcal{I}^n(\Gamma)$ does not have any assignment σ that weakly satisfies all the constraints to obtain the integrality gap. Towards this, we study the sphere colorings $f_n : \mathbb{S}^n \rightarrow \mathcal{B}$ that weakly satisfy all the constraints of $\mathcal{I}^n(\Gamma)$.

We now return to examining 3-LIN in terms of $\mathcal{I}^n(P)$ for P as in (2). As mentioned earlier, a set of three vectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3$ are a P -configuration with respect to \mathbf{v}_0 if they satisfy Equation (3) and Equation (4). Thus, to show that the instance $\mathcal{I}^n(P)$ does not have any assignment satisfying all the constraints, it suffices to show that for some positive integer n , there is no function $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that satisfies the following condition: For all vectors $\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3 \in \mathbb{S}^n$ are mutually orthogonal and are orthogonal to \mathbf{v}_0 , we have

$$f(\mathbf{v}_1) + f(\mathbf{v}_2) + f(\mathbf{v}_3) \in \{-1, +3\}.$$

Since we allow negations of variables in our instance, we also require such a function f to be folded, i.e., $f(-\mathbf{v}) = -f(\mathbf{v})$ for every $\mathbf{v} \in \mathbb{S}^n$. Such a coloring f trivially does not exist: consider a set of three mutually orthogonal vectors that are all orthogonal to \mathbf{v}_0 , $V = (\mathbf{v}_1, \mathbf{v}_2, \mathbf{v}_3)$ and their negations, $V' = (-\mathbf{v}_1, -\mathbf{v}_2, -\mathbf{v}_3)$. Such a set of vectors is guaranteed to exist if $n \geq 4$. Note that both these are valid P -configurations, but at least one of $f(\mathbf{v}_1) + f(\mathbf{v}_2) + f(\mathbf{v}_3)$, $f(-\mathbf{v}_1) + f(-\mathbf{v}_2) + f(-\mathbf{v}_3)$ does not belong to $\{-1, +3\}$, thus completing the proof that there is no assignment satisfying all the constraints of $\mathcal{I}^n(P)$. Hence, $\mathcal{I}^n(P)$ is an integrality gap instance for 3-LIN, i.e., $\text{fiPCSP}(P, P)$, and this implies the existence of a finite integrality gap instance as well. While there is a direct integrality gap instance for the 3-LIN, for an arbitrary promise template Γ , to show the existence of an integrality gap for the basic SDP relaxation, it is more convenient and practical to show the absence of any assignment $f : \mathbb{S}^n \rightarrow \mathcal{B}$ satisfying all the constraints of $\mathcal{I}^n(\Gamma)$ for some n which then implies the existence of a finite integrality gap instance by a compactness argument.

While the P -configurations in the above proof for 3-LIN are easy to study, in general, proving the absence of sphere coloring is challenging. For example, consider the promise template $\Gamma = (P, Q)$ where $P = \text{Ham}_5\{2, 5\}$, $Q = \text{Ham}_5\{1, 2, 3, 4, 5\}$:

$$P = \{\mathbf{x} \in \mathcal{B}^5 : |\{i \in [5] : x_i = +1\}| \in \{2, 5\}\}, \quad Q = \{\mathbf{x} \in \mathcal{B}^5 : \sum_{i=1}^5 x_i \neq -5\}.$$

Here, a set of P -configurations with respect to a vector \mathbf{v}_0 are five unit vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_5\}$ such that every two distinct vectors have an inner product equal to $-\frac{1}{5}$, and $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = 0$ for every $i \in [5]$. The sphere coloring problem is then to show that there exists n such that for any folded $f: \mathbb{S}^n \rightarrow \mathcal{B}$, there exists a set of five vectors in \mathbb{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 in a topic called *sphere Ramsey theory*. In a striking result using tools from combinatorics, linear algebra, and Banach space theory, Matoušek and Rödl [MR95] proved that every set of affinely independent vectors V whose circumradius is smaller than 1 is sphere Ramsey—i.e., for every r , there exists n large enough such that every r -coloring of \mathbb{S}^n must have a monochromatic set U that is congruent to V . This directly answers the above question regarding sphere coloring of $\mathcal{S}^n(P, Q)$ where $P = \text{Ham}_5\{2, 5\}$, $Q = \text{Ham}_5\{1, 2, 3, 4, 5\}$.

For an arbitrary Boolean symmetric promise template $\Gamma = (P, Q)$, to prove Theorem 1.4, we first reduce the problem into a fixed number of templates using the properties of AT and MAJ polymorphisms, following [BG21]. For some of these templates, the result of Matoušek and Rödl [MR95] directly answers the sphere coloring problem associated with them. For others, we need extra work built on the sphere Ramsey result. We highlight one such promise template $\Gamma = (P, Q)$ where $P = \text{Ham}_k\{1, k\}$, $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ for positive integers $k, b: 0 \leq b \leq k$. In the sphere coloring problem associated with this template, we need to show that for any real $\alpha \in [0, 1]$, in any folded $f: \mathbb{S}^n \rightarrow \mathcal{B}$, there are k vectors all of whose pairwise inner products are equal to α , and exactly b of them are assigned $+1$ according to f . We refer to such a set of vectors $S = \{\mathbf{v}_1, \dots, \mathbf{v}_k\} \subseteq \mathbb{S}^n$ as α -configuration if $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \alpha$ for all $i \neq j \in [k]$.

The sphere Ramsey result shows that there is an α -configuration with b vectors being colored $+1$ only when $b = 0$ or $b = k$. To extend to general b , we prove the following connectivity lemma: Between any two arbitrary α -configurations S, T , there is a path U_1, U_2, \dots, U_L of α -configurations with length $L := L(\alpha, k, n)$ where $U_1 = S$, $U_L = T$, and any two consecutive configurations U_i, U_{i+1} differ in at most one element. The sphere Ramsey result shows that there is an α -configuration with all the vectors being colored $+1$, and by negating these vectors, we get an α -configuration with all the vectors being colored -1 . Finally, using the connectivity lemma between these two configurations, we get that for every $b \in \{0, 1, \dots, k\}$, there is an α -configuration where in exactly b vectors are colored $+1$, thereby finishing the proof of the integrality gap. We also remark that while our integrality proofs are non-constructive in general, we get an explicit integrality gap instance for $\text{fiPCSP}(\Gamma)$ inspired by the connectivity lemma.

4 Robust Algorithms

4.1 CMM is a robust algorithm when MAJ is a polymorphism

We restate Theorem 1.3 for the case of MAJ polymorphisms.

Theorem 4.1. *Let $\Gamma = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_l, Q_l)\}$ be a Boolean promise template with $\text{MAJ} \subseteq \text{Pol}(\Gamma)$. For every $\varepsilon > 0$, there is a randomized polynomial time algorithm that given an instance Φ of $\text{fiPCSP}(\Gamma)$ that is promised to have an assignment satisfying $1 - \varepsilon$ fraction of the constraints, finds an assignment to Φ that satisfies $1 - \tilde{O}_\Gamma(\varepsilon^{\frac{1}{3}})$ fraction¹⁰ of the constraints in expectation.*

In the rest of this subsection, we prove Theorem 4.1. Our strategy is to reduce the problem into a special case in which every predicate pair of Γ is of the form $(P, \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\})$, and then solve the robust version of $\text{fiPCSP}(\Gamma)$ using the algorithm of Charikar, Makarychev, and Makarychev [CMM09].

¹⁰We use O_Γ to denote a hidden constant which depends on the specific template Γ .

For ease of notation, we use $O(\cdot)$ instead of $O_\Gamma(\cdot)$ when Γ is clear from the context. We first remove from Γ all predicate pairs (P, Q) for which $Q = \mathcal{B}^k$ for some integer k since these constraints are trivially satisfied by any assignment. To justify why we can remove these predicate pairs in the robust setting, let Φ be our instance of $\text{fiPCSP}(\Gamma)$. Suppose that there are m constraints in Φ , and $m' = \alpha m$ of them use predicates of the form (P, Q) where $P \subseteq Q = \mathcal{B}^k$. We consider the instance Φ' containing $m - m'$ constraints obtained by deleting the constraints that use predicates of the form (P, Q) where $P \subseteq Q = \mathcal{B}^k$. In the instance Φ' , we are promised that there is a solution satisfying $m - m' - \varepsilon m$ constraints, i.e., $1 - \frac{\varepsilon}{1-\alpha}$ fraction of the constraints. We use the algorithm that we will present later in the subsection on the instance Φ' to get an assignment weakly violating at most $\tilde{O}\left(\left(\frac{\varepsilon}{1-\alpha}\right)^{\frac{1}{3}}\right)(m - m')$ constraints. The same assignment weakly violates at most

$$\tilde{O}\left(\left(\frac{\varepsilon}{1-\alpha}\right)^{\frac{1}{3}}\right)(m - m') = \tilde{O}\left(\left(\frac{\varepsilon}{1-\alpha}\right)^{\frac{1}{3}}\right)(1 - \alpha)m \leq \tilde{O}(\varepsilon^{\frac{1}{3}})m$$

constraints in Φ . Thus, it suffices to study Boolean promise templates where no predicate pair is of the form (P, \mathcal{B}^k) .

We further transform the instance into one in which every predicate pair is of the form (P, Q) with $Q = \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\}$.

Lemma 4.2. *Consider a Boolean promise template $\Gamma = \{(P_1, Q_1), \dots, (P_l, Q_l)\}$ where $P_i \subseteq Q_i \subsetneq \mathcal{B}^{k_i}$ for every $i \in [l]$. There exists a promise template $\Gamma' = \{(P'_1, Q'_1), \dots, (P'_l, Q'_l)\}$ and a polynomial time algorithm such that given an instance Φ of $\text{fiPCSP}(\Gamma)$ over a set of variables V , the algorithm outputs an instance Φ' of $\text{fiPCSP}(\Gamma')$ such that the following hold.*

1. (Completeness.) *For every $\varepsilon > 0$, if an assignment $\sigma : V \rightarrow \mathcal{B}$ strongly satisfies $1 - \varepsilon$ fraction of the constraints in Φ , then σ strongly satisfies at least $1 - O_\Gamma(\varepsilon)$ fraction of the constraints in Φ' as well.*
2. (Soundness.) *For every $\varepsilon > 0$, if an assignment $\sigma : V \rightarrow \mathcal{B}$ weakly satisfies $1 - \varepsilon$ fraction of the constraints in Φ' , then σ weakly satisfies at least $1 - O_\Gamma(\varepsilon)$ fraction of constraints in Φ .*
3. *The resulting promise template Γ' satisfies the below two properties:*
 - (a) *For every $i \in [l']$, Q'_i is equal to $\mathcal{B}^{k'_i} \setminus \{(-1, -1, \dots, -1)\}$ for some positive integer k'_i .*
 - (b) *If $\text{MAJ} \subseteq \text{Pol}(\Gamma)$, then, $\text{MAJ} \subseteq \text{Pol}(\Gamma')$.*

Proof. We obtain the above transformation in two steps. First, we construct a Boolean promise template Γ^* from Γ as follows:

$$\Gamma^* := \left\{ \left(P_i, \mathcal{B}^{k_i} \setminus \{\mathbf{b}\} \right) : i \in [l], \mathbf{b} \in \mathcal{B}^{k_i} \setminus Q_i \right\}$$

Note that for every predicate pair $(P, Q) \in \Gamma^*$, there is a predicate pair $(P, Q') \in \Gamma$ with $Q' \subseteq Q$, and thus, $\text{Pol}(\Gamma) \subseteq \text{Pol}(\Gamma^*)$.

Given an instance Φ of $\text{fiPCSP}(\Gamma)$ over a set of variables V , we obtain an instance Φ^* of $\text{fiPCSP}(\Gamma^*)$ over the same set of variables V as follows. We order the constraints of Φ as C_1, C_2, \dots, C_m . Consider a constraint $C_j = (x_{j,1}, x_{j,2}, \dots, x_{j,k_j})$ in Φ using the predicate pair (P_i, Q_i) . In the instance Φ^* , we add $2^{k_i} - |Q_i|$ constraints $C_{j,\mathbf{b}}$ associated with every $\mathbf{b} \in \mathcal{B}^{k_i} \setminus Q_i$. The constraint $C_{j,\mathbf{b}}$ uses the same tuple of literals as C_j but uses the predicate pair $(P_i, \mathcal{B}^{k_i} \setminus \{\mathbf{b}\})$. We analyze the completeness and soundness of this reduction.

1. (Completeness.) If $\sigma : V \rightarrow \mathcal{B}$ strongly satisfies a constraint C_j , then σ strongly satisfies $C_{j,\mathbf{b}}$ for every $\mathbf{b} \in \mathcal{B}^{k_i} \setminus Q_i$. If σ strongly satisfies $(1 - \varepsilon)m$ constraints in Φ , then the number of constraints that σ does not strongly satisfy in Φ^* is at most $2^K \varepsilon m$, where $K := \max_{i \in [l]} k_i$. Thus, σ strongly satisfies at least $m^* - 2^K \varepsilon m \geq (1 - 2^K \varepsilon)m^*$ constraints, where m^* is the number of constraints in Φ^* . Here, we are using the fact that $m^* \geq m$, as we have $Q_i \neq \mathcal{B}^{k_i}$ for every $i \in [k]$.
2. (Soundness.) Suppose that $\sigma : V \rightarrow \mathcal{B}$ weakly satisfies $1 - \varepsilon$ fraction of the constraints in Φ^* . As σ weakly violates at most $\varepsilon m^* \leq 2^K \varepsilon m$ constraints, for at least $1 - 2^K \varepsilon$ fraction of the original constraints C_j , σ weakly satisfies $C_{j,\mathbf{b}}$ for every $\mathbf{b} \in \mathcal{B}^{k_i} \setminus Q_i$. For these constraints C_j , σ weakly satisfies C_j as well, and thus, σ weakly satisfies at least $1 - 2^K \varepsilon$ fraction of the constraints in Φ .

Next, we transform Γ^* to Γ' to ensure that every weak predicate is of the form $\mathcal{B}^k \setminus \{(-1, \dots, -1)\}$. We use the following entry-wise product notation: for a pair of vectors $\mathbf{u}, \mathbf{v} \in \mathbb{R}^k$, we let

$$\mathbf{u} \odot \mathbf{v} := (u_1 v_1, u_2 v_2, \dots, u_k v_k).$$

For a predicate $P \subseteq \mathcal{B}^k$, we let

$$P \odot \mathbf{v} := \{\mathbf{u} \odot \mathbf{v} : \mathbf{u} \in P\}.$$

We let Γ' be the following.

$$\Gamma' := \left\{ \left(P \odot (-\mathbf{b}), \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\} \right) : \left(P, \mathcal{B}^k \setminus \{\mathbf{b}\} \right) \in \Gamma^* \right\}$$

Given the instance Φ^* of $\text{fiPCSP}(\Gamma^*)$ that is obtained from Φ , we construct an instance Φ' of $\text{fiPCSP}(\Gamma')$ on the same set of variables as follows. Consider a constraint $C_{j,\mathbf{b}}$ using the predicate pair $(P_i, \mathcal{B}^{k_i} \setminus \{\mathbf{b}\})$. We have a constraint $C'_{j,\mathbf{b}}$ in Φ' using the predicate pair $(P_i \odot \mathbf{b}, \mathcal{B}^{k_i} \setminus \{(-1, -1, \dots, -1)\})$. However, in the constraint $C'_{j,\mathbf{b}}$, we negate the literals corresponding to the indices $p \in [k_i]$ where $b_p = 1$. More formally, we define the constraint $C'_{j,\mathbf{b}} = (y_{j,1}, y_{j,2}, \dots, y_{j,k_i})$ as follows.

$$y_{j,p} := \begin{cases} x_{j,p}, & \text{if } b_p = -1. \\ \overline{x_{j,p}}, & \text{if } b_p = 1. \end{cases}$$

For the constant literals, $\overline{x_{j,p}}$ refers to a change of sign. An assignment $\sigma : V \rightarrow \mathcal{B}$ strongly (and respectively weakly) satisfies $C'_{j,\mathbf{b}}$ if and only if σ strongly (and respectively weakly) satisfies $C_{j,\mathbf{b}}$. Thus, we get that Φ' satisfies the completeness and soundness properties of the lemma. Finally, the operation of negating a subset of coordinates preserves all polymorphisms that are folded, and thus, if $\text{MAJ} \subseteq \text{Pol}(\Gamma)$, $\text{MAJ} \subseteq \text{Pol}(\Gamma')$ as well. \square

Given an instance Φ of $\text{fiPCSP}(\Gamma)$, we transform it to an instance Φ' of $\text{fiPCSP}(\Gamma')$ using Lemma 4.2. If Φ is promised to have an assignment strongly satisfying at least $1 - \varepsilon$ fraction of the constraints, then Φ' has an assignment strongly satisfying $1 - O(\varepsilon)$ fraction of the constraints as well. If there is a polynomial time robust algorithm that outputs an assignment weakly satisfying $1 - f(\varepsilon)$ fraction of the constraints, then we can use this assignment to obtain a robust algorithm for $\text{fiPCSP}(\Gamma)$ as well. Thus, a polynomial time robust algorithm for $\text{fiPCSP}(\Gamma')$ gives a polynomial time robust algorithm for $\text{fiPCSP}(\Gamma)$ as well.

Given a Boolean promise template Γ in which every predicate pair is of the form $(P, \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\})$ with $\text{MAJ} \subseteq \text{Pol}(\Gamma)$, we show that the robust algorithm of Charikar, Makarychev, and Makarychev [CMM09] for 2-SAT generalizes and gives a robust algorithm for the Boolean folded PCSP Γ as well. First, we state their algorithm.

1. Given an instance Φ of $\text{fiPCSP}(\Gamma)$ containing n variables u_1, u_2, \dots, u_n , solve the Boolean folded idempotent basic SDP and obtain a set of vectors $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_n$ and objective value $\varepsilon > 0$. Let $\boldsymbol{\mu} \in \mathbb{R}^n$ denote the first moments and $\Sigma \in \mathbb{R}^{n \times n}$ be the Gram matrix of these vectors.

$$\begin{aligned}\boldsymbol{\mu}_i &= \langle \mathbf{v}_i, \mathbf{v}_0 \rangle \quad \forall i \in [n] \\ \Sigma_{i,j} &= \langle \mathbf{v}_i, \mathbf{v}_j \rangle \quad \forall i, j \in [n]\end{aligned}$$

2. Sample an n dimensional Gaussian $\zeta \sim \mathcal{N}(\mathbf{0}, \Sigma)$. Note that Σ is positive semidefinite. We can also equivalently sample $\zeta' \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$ and use $\langle \mathbf{v}_i, \zeta' \rangle$ in the place of ζ_i in step 4 below. We stick with $\zeta \sim \mathcal{N}(\mathbf{0}, \Sigma)$ for ease of notation in the analysis.
3. Set¹¹ $\gamma = \varepsilon^{\frac{2}{3}}$.
4. For each $i \in [n]$, round as follows

$$\sigma(u_i) = \begin{cases} +1 & \zeta_i \geq -\mu_i/\gamma. \\ -1 & \text{otherwise.} \end{cases}$$

We shall prove the following guarantee about the algorithm.

Theorem 4.3. *Let $\Gamma = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_l, Q_l)\}$ be a Boolean promise template such that $\text{MAJ} \subseteq \text{Pol}(\Gamma)$ where $P_i \subseteq Q_i = \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\}$ for every $i \in [l]$. Let Φ be an instance of $\text{fiPCSP}(\Gamma)$ over n variables and using m constraints for which the basic SDP relaxation has a solution with error value at most εm . Then, the assignment σ output by the above CMM algorithm weakly satisfies $1 - \tilde{O}_\Gamma(\varepsilon^{1/3})$ fraction of the constraints in expectation.*

Theorem 4.3, together with Lemma 4.2 and Theorem 2.16, completes the proof of Theorem 4.1. Suppose that $\Gamma = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_l, Q_l)\}$ is a Boolean promise template with $\text{MAJ} \subseteq \text{Pol}(\Gamma)$ with $Q_i = \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\}$ for every $i \in [l]$. Given $\varepsilon > 0$ and an instance Φ of $\text{fiPCSP}(\Gamma)$ that is promised to have an assignment satisfying $1 - \varepsilon$ fraction of the constraints, we solve the basic SDP relaxation of the instance. Without loss of generality, we assume that $\varepsilon > \frac{1}{m}$, where m is the number of constraints in the instance. By setting $\delta = \varepsilon$ in Theorem 2.16, we obtain that the basic SDP relaxation has an error of at most $2\varepsilon m$. Finally, Theorem 4.3 implies that the output assignment of the above algorithm weakly satisfies $1 - \tilde{O}_\Gamma(\varepsilon^{1/3})$ of the constraints in expectation.

Towards proving Theorem 4.3, we analyze the probability that the output assignment weakly satisfies each constraint separately. Fix a constraint C_j using the predicate pair (P, Q) of arity k with $P \subseteq Q = \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\}$. Suppose that the basic SDP solution has error equal to c on this constraint, i.e., $\varepsilon_j = c$. Our goal is to upper bound the probability that the rounded solution violates the constraint Q by a function of ε and c . Using the fact that the expected value of c over all the constraints is at most ε , we get our required robustness guarantee. More formally, we prove the following.

Lemma 4.4. *Fix $j \in [m]$, and suppose that the basic SDP solution has an error value equal to c on C_j , i.e.,*

¹¹We change the parameter slightly – in the original CMM algorithm, γ is set to be $\sqrt{\varepsilon}$.

$\varepsilon_j = c$. Then, the probability that σ does not weakly satisfy C_j is at most

$$O\left(\sqrt{\varepsilon} + \sqrt{\left(\gamma\sqrt{\log\frac{1}{\varepsilon}} + 2c\right)\log\frac{1}{\varepsilon} + \frac{c}{\gamma}}\right).$$

By summing over all the constraints, and using linearity of expectation, the expected number of constraints that are not weakly satisfied by σ is at most

$$O\left(\sqrt{\varepsilon}m + \sum_{j \in [m]} \left(\sqrt{\left(\gamma\sqrt{\log\frac{1}{\varepsilon}} + 2\varepsilon_j\right)\log\frac{1}{\varepsilon} + \frac{\varepsilon_j}{\gamma}}\right)\right). \quad (5)$$

As the basic SDP has a total error at most εm , the average value of ε_j over $j \in [m]$ is at most ε . Also note that the expression in Equation (5) is a concave function of ε_j . Thus, using Jensen's inequality, we get that the expected number of constraints that are not weakly satisfied by σ is at most

$$O\left(\sqrt{\varepsilon}m + m \cdot \left(\sqrt{\left(\gamma\sqrt{\log\frac{1}{\varepsilon}} + 2\varepsilon\right)\log\frac{1}{\varepsilon} + \frac{\varepsilon}{\gamma}}\right)\right) \leq \tilde{O}(m\varepsilon^{1/3})$$

This completes the proof of Theorem 4.3. In the rest of the subsection, we prove Lemma 4.4.

For our predicate $P \subseteq \mathcal{B}^k$, let \mathcal{P} be the convex hull of P , where the tuples are viewed as vectors in \mathbb{R}^k .

$$\mathcal{P} := \left\{ \sum_{\mathbf{a} \in P} \lambda_{\mathbf{a}} \mathbf{a} : \lambda_{\mathbf{a}} \geq 0 \ \forall \mathbf{a} \in P, \sum_{\mathbf{a} \in P} \lambda_{\mathbf{a}} = 1 \right\}.$$

We prove the following property about \mathcal{P} using the fact that $\text{MAJ} \subseteq \text{Pol}(P, \mathcal{B}^k \setminus \{(-1, \dots, -1)\})$. We recall that $O_{\text{MAJ}}(P)$ denotes the set $\bigcup_{\mathbf{x}_1, \dots, \mathbf{x}_L \in P, L \in \mathbb{N}, \text{odd}} \text{MAJ}_L(\mathbf{x}_1, \dots, \mathbf{x}_L)$ for a predicate $P \subseteq \mathcal{B}^k$.

Lemma 4.5. *Let $P \subseteq \mathcal{B}^k$ be a predicate such that $(-1, -1, \dots, -1) \notin O_{\text{MAJ}}(P)$. Then, there is a hyperplane separating \mathcal{P} from the origin: there exists $\mathbf{w} \in \mathbb{R}^k$, $\mathbf{w} \geq 0$ and $\|\mathbf{w}\|_1 = 1$ such that for every $\mathbf{a} \in \mathcal{P}$, $\langle \mathbf{a}, \mathbf{w} \rangle \geq 0$.*

Proof. Consider the following linear program,

$$\begin{aligned} & \mathbf{maximize:} \quad \eta \\ & \mathbf{subject\ to:} \quad \sum_{i=1}^k w_i = 1 \\ & \quad \forall \mathbf{a} \in P, \eta - \sum_{i=1}^k a_i w_i \leq 0 \\ & \quad \mathbf{w} \geq 0 \end{aligned}$$

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

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

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

$$\sum_{a \in P} a_i \lambda_a < 0.$$

Let N be the least common denominator of the λ_a s. Consider the set of satisfying assignments to P where we take $2N\lambda_a$ copies of a for each $a \in P$. We also add an arbitrary element b of a to our set. As $\sum_{a \in P} a_i \lambda_i N < 0$ for every $i \in [k]$, and $\sum_{a \in P} a_i \lambda_i N$ is an integer, we get that for every $i \in [k]$, $2\sum_{a \in P} a_i \lambda_a N + b \leq 0$. In other words, when we apply MAJ_{2N+1} coordinatewise to this set of assignments in P , we get $(-1, -1, \dots, -1)$. As (P, Q) contains Majority of all odd arities as polymorphisms, this implies that the resulting output $(-1, -1, \dots, -1)$ is in Q , a contradiction.

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

Now we use this lemma to complete the proof of Lemma 4.4. Recall that our goal is to lower bound the probability that the assignment σ output by the above algorithm weakly satisfies the constraint C_j . Suppose that $C_j = (x_{j,1}, x_{j,2}, \dots, x_{j,k})$, and uses the predicate pair (P, Q) where $P \subseteq Q = \mathcal{B}^k \setminus \{(-1, -1, \dots, -1)\}$. Let S_j denote the set of variables appearing in C_j . For each $l \in [k]$, let $\hat{\mathbf{v}}_l := \mathbf{v}(x_{j,l})$, where the function \mathbf{v} is defined as in (1). Let $\mu^{(j)}$ and $\Sigma^{(j)}$ be defined as follows:

$$\begin{aligned} \mu_l^{(j)} &:= \langle \hat{\mathbf{v}}_l, \mathbf{v}_0 \rangle & \forall l \in [k], \\ \Sigma_{l,l'}^{(j)} &:= \langle \hat{\mathbf{v}}_l, \hat{\mathbf{v}}_{l'} \rangle & \forall l, l' \in [k]. \end{aligned}$$

Since $\varepsilon_j = c$, we have that

$$\sum_{\substack{f: S_j \rightarrow \mathcal{B} \\ f(C_j) \in P}} \lambda_j(f) = 1 - c.$$

Using Lemma 4.5, we get $\mathbf{w} \in \mathbb{R}^k$ with $\mathbf{w} \geq 0$ and $\|\mathbf{w}\|_1 = 1$ such that $\langle \mathbf{w}, \mathbf{a}_i \rangle \geq 0$ for all $\mathbf{a}_i \in P$. Combining this with the above properties of the basic SDP solution, we get the following.

1. (First moment). We have

$$\begin{aligned} \langle \mathbf{w}, \mu^{(j)} \rangle &= \sum_{f: S_j \rightarrow \mathcal{B}} \lambda_j(f) \langle \mathbf{w}, f(C_j) \rangle \\ &\geq -c \text{ (Using Lemma 4.5 and } -1 \leq \langle \mathbf{w}, \mathbf{a} \rangle \leq 1 \forall \mathbf{a} \in \mathcal{B}^k \text{)} \end{aligned}$$

2. (Second moment). We have

$$\begin{aligned} \mathbf{w}^T \Sigma^{(j)} \mathbf{w} &= \sum_{f: S_j \rightarrow \mathcal{B}} \lambda_j(f) \langle \mathbf{w}, f(C_j) \rangle^2 \\ &\leq \sum_{\substack{f: S_j \rightarrow \mathcal{B} \\ f(C_j) \in P}} \lambda_j(f) (\langle \mathbf{w}, f(C_j) \rangle)^2 + c \\ &\leq \sum_{\substack{f: S_j \rightarrow \mathcal{B} \\ f(C_j) \in P}} \lambda_j(f) \langle \mathbf{w}, f(C_j) \rangle + c \leq \langle \mathbf{w}, \mu \rangle + 2c. \end{aligned}$$

We do casework on the value of $\langle \mathbf{w}, \mu^{(j)} \rangle$. First, consider the case that $\langle \mathbf{w}, \mu^{(j)} \rangle \geq \kappa = \gamma \sqrt{\log \frac{1}{\varepsilon}}$. As $\|\mathbf{w}\|_1 = 1$, and $\mathbf{w} \geq 0$, there exists $i \in [k]$ such that $\mu_i^{(j)} \geq \kappa$. As $\zeta_i \sim \mathcal{N}(0, 1)$, using Proposition 2.21, with probability at least $1 - \sqrt{\varepsilon}$, we have $\zeta_i \geq -\frac{\mu_i^{(j)}}{\gamma}$. Thus, with probability at least $1 - \sqrt{\varepsilon}$, the rounded solution satisfies Q .

Henceforth, we assume $\mathbf{w}^T \boldsymbol{\mu}^{(j)} < \kappa$. For notational convenience let $\mathbf{t} = -\boldsymbol{\mu}^{(j)}/\gamma$. We have

$$\mathbf{w}^T \mathbf{t} \leq \frac{c}{\gamma} \quad (6)$$

and $\mathbf{w}^T \Sigma^{(j)} \mathbf{w} \leq \kappa + 2c$. From Proposition 2.22, $\langle \mathbf{w}, \boldsymbol{\zeta} \rangle \sim \mathcal{N}(0, \mathbf{w}^T \Sigma^{(j)} \mathbf{w})$. Thus, using Proposition 2.21, with probability at least $1 - \sqrt{\varepsilon}$, we have that

$$|\mathbf{w}^T \boldsymbol{\zeta}| \leq O\left(\sqrt{(\kappa + 2c) \log \frac{1}{\varepsilon}}\right) \quad (7)$$

Note that the rounded solution does not satisfy Q only if $\mathbf{t} \geq \boldsymbol{\zeta}$. We now upper bound the probability that this occurs. Together with Equation (6) and Equation (7), $\mathbf{t} \geq \boldsymbol{\zeta}$ implies that

$$0 \leq \mathbf{w}^T (\mathbf{t} - \boldsymbol{\zeta}) \leq O\left(\sqrt{(\kappa + 2c) \log \frac{1}{\varepsilon}} + \frac{c}{\gamma}\right).$$

Take some coordinate with $w_i \geq 1/k$ and note that

$$t_i - \zeta_i \in \left[0, O\left(\sqrt{(\kappa + 2c) \log \frac{1}{\varepsilon}} + \frac{c}{\gamma}\right)\right].$$

Since t_i is a constant, this occurs with probability

$$O\left(\sqrt{(\kappa + 2c) \log \frac{1}{\varepsilon}} + \frac{c}{\gamma}\right),$$

using Proposition 2.20 applied to $\zeta_i \sim \mathcal{N}(0, 1)$. Thus, the probability that the assignment σ does not satisfy Q is at most

$$O\left(\sqrt{\varepsilon} + \sqrt{(\kappa + 2c) \log \frac{1}{\varepsilon}} + \frac{c}{\gamma}\right).$$

This completes the proof of Lemma 4.4.

4.2 Warm-up for AT: Random threshold rounding algorithm for Dual-Horn SAT

As a stepping-stone for our algorithm for AT presented in the next section, we present a robust algorithm for the Dual-Horn SAT CSP. The formal template Dual-Horn SAT allows for constraints of arbitrarily large arity, but we consider a finite fragment. In particular, our template $\Gamma = \{P_1, P_2, \dots, P_l\}$ has $P_i = \mathcal{B}^{k_i} \setminus \{(-1, -1, \dots, -1)\}$ for $i \in [l]$. Specifically to emulate Dual-Horn SAT, we restrict ourselves to instances of $\text{fiPCSP}(\Gamma)$ such that each clause has at most one negated literal and no constant literals appear.

A robust algorithm for Dual-Horn SAT was found previously by Zwick [Zwi98], and a matching hardness result was obtained by Guruswami and Zhou [GZ12]. We now present a robust algorithm for Dual-Horn SAT achieving similar guarantees as [Zwi98]. The random thresholding idea that we use here serves as a good warmup for the AT algorithm presented in the next subsection.

1. Given an instance Φ of $\text{fiPCSP}(\Gamma)$ containing n variables $\{u_1, u_2, \dots, u_n\}$, solve the basic SDP and obtain a set of vectors $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_n$ and objective value $\varepsilon > 0$. Let $\mu \in \mathbb{R}^n$ denote the first moments.

$$\mu_i = \langle \mathbf{v}_i, \mathbf{v}_0 \rangle \quad \forall i \in [n]$$

2. Let T be a geometric progression with first term $\sqrt{\varepsilon}$, last term $1/K$ and spacing between terms is at least K , where K is the maximum arity of the predicates in Γ .
3. Sample a uniformly random threshold $t \in T$.
4. For each $i \in [n]$, round as follows

$$\sigma(x_i) = \begin{cases} +1 & \text{if } \mu_i \geq t - 1. \\ -1 & \text{otherwise.} \end{cases}$$

Theorem 4.6. *Let Φ be an instance of $\text{fiPCSP}(\Gamma)$ such that each clause has at most one negated literal and no constant literals for which there is a basic SDP solution with error at most ε . Then, the assignment σ output by the above algorithm satisfies $1 - O_\Gamma(1/\log(1/\varepsilon))$ fraction of the constraints of Φ in expectation.*

Proof. As with the analysis of the MAJ polymorphism, we fix a single constraint C_j using the predicate P and analyze the probability that σ satisfies C_j . Since the average SDP error over the constraints is equal to ε , by Markov's inequality, for at least $1 - \sqrt{\varepsilon}$ fraction of the constraints, the SDP error value is at most $\sqrt{\varepsilon}$. Henceforth, we assume that $c \leq \sqrt{\varepsilon}$.

Suppose that the constraint C_j uses the literals x_1, \dots, x_k with some potentially equal. As an abuse of notation, we let μ_i denote the inner product of \mathbf{v}_0 with $\mathbf{v}(x_i)$. First, consider the case when none of the variables are negated. Using the fact that the local assignments in the SDP solution have weight at least $1 - \sqrt{\varepsilon}$ fraction on assignments in P , we get that

$$\sum_{i=1}^k \mu_i \geq (1 - \sqrt{\varepsilon})(-(k-2)) + \sqrt{\varepsilon}(-k) = -(k-2) - 2\sqrt{\varepsilon}$$

By the pigeonhole principle, we must have some i with

$$\mu_i \geq \frac{-(k-2) - 2\sqrt{\varepsilon}}{k} = -1 + \frac{2 - 2\sqrt{\varepsilon}}{k} \geq -1 + \frac{1}{K}$$

where in the final step, we assumed that ε is small enough that $2\sqrt{\varepsilon} \leq 1$. Thus, $\mu_i \geq t - 1$ for all $t \in T$. In this case, $\sigma(x_i) = +1$, and the constraint is satisfied by σ .

Next, consider the case in which exactly one literal is a negated variable in the constraint. Without loss of generality, assume that x_1 is negated. Then, we have that

$$-\mu_1 + \mu_2 + \dots + \mu_k \geq (1 - \sqrt{\varepsilon})(-(k-2)) + \sqrt{\varepsilon}(-k) = -(k-2) - 2\sqrt{\varepsilon}.$$

In particular, there exists an integer $i \in \{2, 3, \dots, k\}$ such that

$$\mu_i \geq \frac{\mu_1 - (k-2) - 2\sqrt{\varepsilon}}{(k-1)} = -1 + \frac{\mu_1 + 1 - 2\sqrt{\varepsilon}}{k-1}$$

The algorithm rounds x_i to -1 only if $t > \frac{\mu_1+1-2\sqrt{\varepsilon}}{k-1}$. On the other hand, the algorithm rounds x_1 to $+1$ only if $t \leq \mu_1 + 1$. Since there is at most one element in T in $(\frac{\mu_1+1-2\sqrt{\varepsilon}}{k-1}, \mu_1 + 1)$, there exists at most one $t \in T$ which would round x_i to -1 but x_1 to $+1$. Therefore, the probability the constraint is satisfied is at least $1 - 1/|T| = 1 - O_\Gamma(1/\log(1/\varepsilon))$. By linearity of expectation over all the constraints with error at most $\sqrt{\varepsilon}$, we get the required claim. \square

Theorem 4.6 implies a polynomial time robust algorithm for Dual-Horn SAT: Suppose that the template $\Gamma = \{P_1, P_2, \dots, P_l\}$ has $P_i = \mathcal{B}^{k_i} \setminus \{(-1, -1, \dots, -1)\}$ for $i \in [l]$. Consider an instance Φ of fiPCSP(Γ) such that each clause has at most one negated literal and no constant literals appear, with the property that there is an assignment that strongly satisfies $(1 - \varepsilon)$ fraction of the constraints. Without loss of generality, we can assume that $\varepsilon \geq \frac{1}{m}$, where m is the number of constraints in Φ . By setting $\delta = \sqrt{\varepsilon} - \varepsilon$ in Theorem 2.16, we can obtain a solution to basic SDP relaxation of Φ with error at most $\sqrt{\varepsilon}m$. Theorem 4.6 then implies that the output of the above algorithm satisfies $1 - O_\Gamma(1/\log(1/\varepsilon))$ fraction of the constraints of Φ in expectation.

Remark 4.7. We remark that the above algorithm applies directly to Horn-SAT as well. Furthermore, for an arbitrary Boolean promise template $\Gamma = \{(P_1, Q_1), (P_2, Q_2), \dots, (P_l, Q_l)\}$ with OR polymorphisms¹² $\text{OR} \subseteq \text{Pol}(\Gamma)$, it's been proved [BG21] that $\text{OR} \subseteq \text{Pol}(\Gamma')$ where Γ' is the CSP containing the predicates $\{Q_1, Q_2, \dots, Q_l\}$. By Schaefer's theorem [Sch78], any template Γ with $\text{OR} \subseteq \text{Pol}(\Gamma)$ is ppp-definable from k -Horn-SAT, and thus, our algorithm proves that there is a polynomial time robust algorithm for PCSP(Γ) (i.e., negated literals are not allowed) for every Boolean promise template Γ with $\text{OR} \subseteq \text{Pol}(\Gamma)$.

4.3 Algorithm for AT

We now show how to combine the ideas in the previous two subsections to obtain a polynomial time robust algorithm for every Boolean promise template $\Gamma = \{(P_1, Q_1), \dots, (P_l, Q_l)\}$, $P_i \subseteq Q_i \subseteq \mathcal{B}^{k_i}$ with $\text{AT} \subseteq \text{Pol}(\Gamma)$. Similar to the MAJ polymorphisms case, we can without loss of generality assume that $Q_i \neq \mathcal{B}^{k_i}$ for every $i \in [l]$. We first reduce an arbitrary promise template Γ with $\text{AT} \subseteq \text{Pol}(\Gamma)$ to a specific template that we will work with. For a vector $\mathbf{w} \in \mathbb{R}^k$, define $\text{sgn}(\mathbf{w})_i$ to be -1 if $w_i < 0$ and $+1$ otherwise. Define Γ_{AT} to be the following infinite family of weighted hyperplane predicates:

$$\Gamma_{\text{AT}} := \{(P_{\mathbf{w},b} := \{\mathbf{x} \in \mathcal{B}^k : \langle \mathbf{w}, \mathbf{x} \rangle = b\}, Q_{\mathbf{w},b} := \mathcal{B}^k \setminus \{\text{sgn}(\mathbf{w}), -\text{sgn}(\mathbf{w})\}) : \\ k \in \mathbb{N}, b \in \mathbb{Q}, \mathbf{w} \in \mathbb{Q}^k \setminus \mathbf{0}, \mathbf{w} \cdot \text{sgn}(\mathbf{w}) > b, -\mathbf{w} \cdot \text{sgn}(\mathbf{w}) < b\}$$

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

Claim 4.8. $\text{AT} \subseteq \text{Pol}(\Gamma_{\text{AT}})$

Proof. Fix $b \in \mathbb{Q}$ and $\mathbf{w} \in \mathbb{Q}^k \setminus \mathbf{0}$. Let $(P_{\mathbf{w},b}, Q_{\mathbf{w},b})$ be the corresponding predicate for these values. It suffices to show that $\text{AT} \subseteq \text{Pol}(P_{\mathbf{w},b}, Q_{\mathbf{w},b})$. Fix an odd arity L and pick $\mathbf{x}_1, \dots, \mathbf{x}_L \in P_{\mathbf{w},b}$. Observe that

$$\text{AT}(\mathbf{x}_1, \dots, \mathbf{x}_L) = \text{sgn}(\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L).$$

Further, $\langle \mathbf{w}, (\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L) \rangle = b$. This implies that $\text{sgn}(\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L) \neq \text{sgn}(\mathbf{w})$ as otherwise,

$$b = \langle \mathbf{w}, (\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L) \rangle \geq \langle \mathbf{w}, \text{sgn}(\mathbf{w}) \rangle > b.$$

where we used the fact that the absolute value of each entry in $\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L$ is at least 1. By a similar argument, $\text{sgn}(\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L) \neq -\text{sgn}(\mathbf{w})$. Thus, $\text{AT}(\mathbf{x}_1, \dots, \mathbf{x}_L) \in Q_{\mathbf{w},b}$, as desired. \square

¹² $\text{OR}(x_1, x_2, \dots, x_L)$ is equal to $+1$ if there is $i \in [L]$ with $x_i = +1$, and -1 otherwise.

Let Γ_{const} be the Boolean promise template $\{(\{-1\}, \{-1\}), (\{+1\}, \{+1\})\}$. Since AT is an idempotent family of polymorphisms, we have that $AT \subseteq \text{Pol}(\Gamma_{\text{const}})$. We show that an arbitrary Boolean promise template with AT polymorphisms can be reduced to a finite fragment of $\Gamma_{\text{AT}} \cup \Gamma_{\text{const}}$.

Lemma 4.9. *Let Γ be any Boolean promise template with $\Gamma = \{(P_1, Q_1), \dots, (P_l, Q_l)\}, P_i \subseteq Q_i \subsetneq \mathcal{B}^{k_i}$ with $AT \subseteq \text{Pol}(\Gamma)$. Then, Γ is ppp-definable from a finite Boolean promise template Γ' with $\Gamma' \subseteq \Gamma_{\text{AT}} \cup \Gamma_{\text{const}}$.*

We prove Lemma 4.9 in Section 4.4. Recall that if a promise template Γ is ppp-definable from another promise template Γ' and if $\text{fiPCSP}(\Gamma')$ has a polynomial time robust algorithm, then $\text{fiPCSP}(\Gamma)$ has a polynomial time robust algorithm as well (up to losing constant factors in the error parameter). In the rest of this section, we prove that for any finite Boolean promise template Γ with $\Gamma \subseteq \Gamma_{\text{AT}}$, $\text{fiPCSP}(\Gamma)$ has a robust algorithm, thereby obtaining a robust algorithm for $\text{fiPCSP}(\Gamma)$ for every Boolean promise template Γ with $AT \subseteq \text{Pol}(\Gamma)$. We state our algorithm below.

1. Given an instance Φ of $\text{fiPCSP}(\Gamma)$ containing n variables $V = \{u_1, u_2, \dots, u_n\}$, solve the basic SDP and obtain a set of vectors $\mathbf{v}_0, \mathbf{v}_1, \dots, \mathbf{v}_n$ and objective value $\varepsilon > 0$.
2. Sample an n dimensional Gaussian $\zeta \sim \mathcal{N}(\mathbf{0}, \mathbf{I})$.
3. Choose δ uniformly at random from $\{p, rp, \dots, r^\kappa p\}$ where $p = \varepsilon^{c_1}$, and $r = \kappa = \Theta\left(\frac{\log \frac{1}{\varepsilon}}{\log \log \frac{1}{\varepsilon}}\right)$ such that $r^\kappa p = \varepsilon^{c_2}$. Here, c_1 and c_2 are two arbitrary real constants satisfying $0 < c_2 < c_1 < 0.25$.
4. For every $i \in [n]$, let $\mathbf{v}_i = \alpha_i \mathbf{v}_0 + \mathbf{v}'_i$, where \mathbf{v}'_i is orthogonal to \mathbf{v}_0 . We output an assignment σ as follows.

$$\sigma(u_i) = \begin{cases} -1, & \text{if } \langle \zeta, \mathbf{v}'_i \rangle \geq \delta \alpha_i |\langle \zeta, \mathbf{v}_0 \rangle|. \\ +1, & \text{otherwise.} \end{cases}$$

We now analyze the algorithm.

Theorem 4.10. *Let Γ be a finite Boolean promise template such that $\Gamma \subseteq \Gamma_{\text{AT}} \subseteq \Gamma_{\text{const}}$. Let Φ be an instance of $\text{fiPCSP}(\Gamma)$ for which there is a basic SDP solution with average error at most $\varepsilon > 0$. Then, the assignment output by the above algorithm satisfies at least $1 - O_\Gamma\left(\frac{\log \log \frac{1}{\varepsilon}}{\log \frac{1}{\varepsilon}}\right)$ fraction of constraints of Φ in expectation.*

Theorem 4.10 together with Lemma 4.9 implies a polynomial time robust algorithm for Boolean promise templates with AT polymorphisms. Consider an instance Φ of $\text{fiPCSP}(\Gamma)$ with $\Gamma \subseteq \Gamma_{\text{AT}} \cup \Gamma_{\text{const}}$ with the property that there is an assignment that strongly satisfies $(1 - \varepsilon)$ fraction of the constraints. Without loss of generality, we can assume that $\varepsilon \geq \frac{1}{m}$, where m is the number of constraints in Φ . By setting $\delta = \sqrt{\varepsilon} - \varepsilon$ in Theorem 2.16, we can obtain a solution to basic SDP relaxation of Φ with error at most $\sqrt{\varepsilon}m$. Theorem 4.10 then implies that the output of the above algorithm satisfies $1 - O_\Gamma\left(\frac{\log \log \frac{1}{\varepsilon}}{\log \frac{1}{\varepsilon}}\right)$ fraction of the constraints of Φ in expectation. Together with Lemma 4.4, this completes the proof of Theorem 1.3.

For ease of notation, we just use $O()$ instead of $O_\Gamma()$ when Γ is clear from the context. As before, we prove Theorem 4.10 by proving a lower bound on the probability that a particular constraint is satisfied. Consider the constraint $C = (x_1, \dots, x_k)$ using the predicate pair (P, Q) where $P \subseteq Q \subseteq \mathcal{B}^k$, and let c denote the error of the SDP solution on constraint C_j . As the average value of c over all the constraints is at most ε , using Markov's inequality, at least $1 - \sqrt{\varepsilon}$ fraction of the constraints have SDP error at most $\sqrt{\varepsilon}$. We restrict our analysis to these constraints with SDP error $c \leq \sqrt{\varepsilon}$ and show that the rounded solution violates

the predicate Q with probability at most $O\left(\frac{\log \log \frac{1}{\varepsilon}}{\log \frac{1}{\varepsilon}}\right)$, thereby proving Theorem 4.10. Finally, we note that the rounding algorithm preserves the properties of idempotence and folding: specifically, $\pm \mathbf{v}_0$ are rounded to ± 1 , respectively. Moreover, for any vector \mathbf{v} , the algorithm ensures that \mathbf{v} and $-\mathbf{v}$ are rounded to distinct values.

We break the analysis into cases based on the predicate pair (P, Q) . We first consider the case when $P = Q$ and $P \subseteq \text{Ham}_k\{0, k\}$, which is a generalization of Γ_{const} .

Lemma 4.11. *Let $\Gamma = (P, Q)$ be a finite Boolean promise template with $\text{AT} \subseteq \text{Pol}(\Gamma)$ and $P = Q$, $P \subseteq \text{Ham}_k\{0, k\}$ for an integer $k \geq 1$. Let Φ be an instance of $\text{fiPCSP}(\Gamma)$, and in the basic SDP solution of Φ , suppose that the constraint C_j using the predicate pair (P, Q) has error at most $\sqrt{\varepsilon}$. Then, the assignment σ output by the above algorithm does not weakly satisfy C_j with probability at most $O\left(\frac{\log \log \frac{1}{\varepsilon}}{\log \frac{1}{\varepsilon}}\right)$.*

We defer the proof of Lemma 4.11 to Appendix B.

We now consider the case when $P = P_{\mathbf{w}, b}$ and $Q = Q_{\mathbf{w}, b}$ for some $k \in \mathbb{N}, b \in \mathbb{Q}, \mathbf{w} \in \mathbb{Q}^k \setminus \mathbf{0}, \mathbf{w} \cdot \text{sgn}(\mathbf{w}) > b, -\mathbf{w} \cdot \text{sgn}(\mathbf{w}) < b$. For ease of notation, we let x_1, x_2, \dots, x_k denote the literals assigned to C_j , with $\mathbf{v}_1, \dots, \mathbf{v}_k$ the corresponding values of $\mathbf{v}(x_1), \dots, \mathbf{v}(x_k)$. We let S denote the set of variables in the scope of C_j . We have that

$$\sum_{\substack{f: S \rightarrow \mathcal{B} \\ f(C_j) \in P}} \lambda_j(f) \geq 1 - \sqrt{\varepsilon}.$$

Let $\mathbf{v}_s = \sum_{i \in [k]} \mathbf{w}_i \mathbf{v}_i$, and let the component of \mathbf{v}_s along \mathbf{v}_0 be $\alpha \mathbf{v}_0$, and the component normal to \mathbf{v}_0 be equal to \mathbf{v}'_s .

$$\mathbf{v}_s = \alpha \mathbf{v}_0 + \mathbf{v}'_s, \langle \mathbf{v}_0, \mathbf{v}'_s \rangle = 0.$$

We use the first and second moments properties of the vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k$ to get the following.

1. (First moments). We have

$$\begin{aligned} \alpha &= \sum_{f: S \rightarrow \mathcal{B}} \lambda_j(f) \langle \mathbf{w}, f(C_j) \rangle \\ &= \sum_{\substack{f: S \rightarrow \mathcal{B} \\ f(C_j) \in P}} \lambda_j(f) \langle \mathbf{w}, f(C_j) \rangle + \sum_{\substack{f: S \rightarrow \mathcal{B} \\ f(C_j) \notin P}} \lambda_j(f) \langle \mathbf{w}, f(C_j) \rangle \\ &= \sum_{\substack{f: S \rightarrow \mathcal{B} \\ f(C_j) \in P}} \lambda_j(f) b + \sum_{\substack{f: S \rightarrow \mathcal{B} \\ f(C_j) \notin P}} \lambda_j(f) \langle \mathbf{w}, f(C_j) \rangle \\ &= b + \beta \end{aligned}$$

where

$$\beta = \sum_{\substack{f: S \rightarrow \mathcal{B} \\ f(C_j) \notin P}} \lambda_j(f) (b - \langle \mathbf{w}, f(C_j) \rangle)$$

We have $|\beta| = O(\sqrt{\varepsilon})$.

2. (Second moments). We have

$$\|\mathbf{v}_s\|_2^2 = \sum_{i, j \in [k]} \mathbf{w}_i \mathbf{w}_j \langle \mathbf{v}_i, \mathbf{v}_j \rangle = \sum_{f: S \rightarrow \mathcal{B}} \lambda_j(f) (\langle f(C_j), \mathbf{w} \rangle)^2 = b^2 + \beta'$$

where $|\beta'| = O(\sqrt{\varepsilon})$.

Thus, we get $\|\mathbf{v}'_s\|_2^2 = \|\mathbf{v}_s\|_2^2 - \alpha^2 = (b^2 + \beta') - (b + \beta)^2$ which is at most $O(\sqrt{\varepsilon})$.

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

Case 1. Suppose that there exists $i \in [k]$ such that $\|\mathbf{v}'_i\|_2 > k\delta r^2$. We claim that in this case, the rounded solution satisfies Q with probability at least $1 - O(\frac{1}{r})$.

From Proposition 2.22, $\langle \zeta, \mathbf{v}'_j \rangle \sim \mathcal{N}(0, \|\mathbf{v}'_j\|_2^2)$ for every $j \in [k]$. Since $\|\mathbf{v}'_i\|_2 > k\delta r^2$, using Proposition 2.20, with probability at least $1 - O(\frac{1}{r})$, we have that $|\langle \zeta, \mathbf{v}'_i \rangle| > k\delta r$.

Recall that $\|\mathbf{v}'_s\|_2 \leq \varepsilon^{0.25} \leq \delta$. Thus, $|\langle \zeta, \mathbf{v}'_s \rangle| \leq r\delta$ with probability at least $O(\frac{1}{r})$. As $|\langle \zeta, \mathbf{v}'_i \rangle| > k\delta r$, there exists $i' \in [k], i' \neq i$, such that $|\langle \zeta, \mathbf{v}'_{i'} \rangle| > \delta r$, and $\langle \zeta, \mathbf{v}'_i \rangle$ and $\langle \zeta, \mathbf{v}'_{i'} \rangle$ have opposite signs.

Furthermore, as $\langle \zeta, \mathbf{v}_0 \rangle \sim \mathcal{N}(0, 1)$, using Proposition 2.21, we get that $|\langle \zeta, \mathbf{v}_0 \rangle| \leq r$ with probability at least $1 - O(\varepsilon)$. Combining this with the above, we conclude that with probability $1 - O(\frac{1}{r})$, all the following hold:

- (1) $|\langle \zeta, \mathbf{v}_0 \rangle| \leq r$.
- (2) $|\langle \zeta, \mathbf{v}'_j \rangle| > \delta r$ for all $j \in \{i, i'\}$.
- (3) $\langle \zeta, \mathbf{v}'_i \rangle$ and $\langle \zeta, \mathbf{v}'_{i'} \rangle$ have opposite signs.

Note that for every $j \in [k]$, $|\alpha_j| = |\langle \mathbf{v}_0, \mathbf{v}_j \rangle| \leq 1$. Using (1), we have $|\delta\alpha_j \langle \zeta, \mathbf{v}_0 \rangle| \leq r\delta$. Thus, from (2), for all $j \in \{i, i'\}$, we deduce that x_j is set to -1 if $\langle \zeta, \mathbf{v}'_j \rangle > 0$, and $+1$ otherwise.

Combining with (3), with probability at least $1 - O(\frac{1}{r})$, i and i' are rounded to different values, which implies that the rounded solution satisfies Q .

Case 2. Suppose that for every $i \in [k]$, we have $\|\mathbf{v}'_i\|_2 \leq \frac{\delta}{2r^2}$.

As $\langle \zeta, \mathbf{v}'_i \rangle \sim \mathcal{N}(0, \|\mathbf{v}'_i\|_2^2)$, using Proposition 2.21, we get that with probability at least $1 - O(\frac{1}{r})$, for every $i \in [k]$, $|\langle \zeta, \mathbf{v}'_i \rangle| \leq \frac{\delta}{2r}$. On the other hand, using Proposition 2.20, we have that $|\langle \zeta, \mathbf{v}_0 \rangle| \geq \frac{1}{r}$ with probability at least $1 - \frac{1}{r}$. Furthermore, As $\alpha_i^2 + \|\mathbf{v}'_i\|_2^2 = 1$ for every $i \in [k]$, we get that $|\alpha_i| \geq 1 - \delta \geq \frac{1}{2}$ for every $i \in [k]$. Thus, with probability at least $1 - O(\frac{1}{r})$, for every $i \in [k]$, x_i is set to be -1 if $\alpha_i \leq 0$, and $+1$ otherwise. Combining this with the fact that $\sum_i \mathbf{w}_i \alpha_i = b + O(\sqrt{\varepsilon})$, and that $\mathbf{w} \cdot \text{sgn}(\mathbf{w}) > b, -\mathbf{w} \cdot \text{sgn}(\mathbf{w}) < b$, for small enough ε , we get the rounded solution has variables assigned $+1$ and -1 . Here, we assumed that ε is at most $c := c(\Gamma)$ for some function c . This is without loss of generality, since when ε is larger than c , a uniformly random assignment satisfies the robust algorithm requirement.

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

$$\frac{\delta}{2r^2} < \|\mathbf{v}'_i\|_2 \leq k\delta r^2$$

Or equivalently,

$$\frac{\|\mathbf{v}'_i\|_2}{kr^2} \leq \delta < \|\mathbf{v}'_i\|_2 2r^2$$

This holds with probability at most $O(\frac{1}{r})$ for every value of $\|\mathbf{v}'_i\|_2$ as we are picking δ from $\{p, rp, \dots, r^k p\}$ uniformly at random with $\kappa = r$.

4.4 Proof of Lemma 4.9

For a predicate $P \subseteq \mathcal{B}^k$, we let $\text{Aff}(P) \subseteq \mathbb{R}^k$ to denote the affine hull of P .

$$\text{Aff}(P) := \left\{ \sum_{\mathbf{a} \in P} \lambda_{\mathbf{a}} \mathbf{a} : \sum_{\mathbf{a} \in P} \lambda_{\mathbf{a}} = 1, \lambda_{\mathbf{a}} \in \mathbb{R} \right\}.$$

We recall that $O_{\text{AT}}(P)$ denotes the set $\bigcup_{\mathbf{x}_1, \dots, \mathbf{x}_L \in P, L \in \mathbb{N}, \text{odd}} \text{AT}_L(\mathbf{x}_1, \dots, \mathbf{x}_L)$ for a predicate $P \subseteq \mathcal{B}^k$. In order to prove Lemma 4.9, we need to use the following lemma implicit in [BG21].

Lemma 4.12. *Let $P \subseteq \mathcal{B}^k$ be a predicate such that there is non-trivial dependence in each coordinate, i.e., for every $i \in [k]$, there exist vectors $\mathbf{x}, \mathbf{y} \in P$ with $x_i = -1$ and $y_i = +1$. Then, $O_{\text{AT}}(P) = \{\text{sgn}(\mathbf{x} - \mathbf{y}) : \mathbf{x}, \mathbf{y} \in \text{Aff}(P), \forall i, x_i \neq y_i\}$.*

We present the proof in Appendix B for the sake of completeness. We also need the following claim.

Claim 4.13. *Let H be a vector space in \mathbb{R}^k such that there is no $\mathbf{y} \in H$ with $y_i > 0$ for all i . Then, there exists $\mathbf{w} \neq 0$ with $w_i \geq 0$ for all i and $\langle \mathbf{w}, \mathbf{y} \rangle = 0$ for all $\mathbf{y} \in H$.*

Proof. Since H and the positive orthant $\mathbb{R}_+^k := \{\mathbf{y} \in \mathbb{R}^k : y_i > 0 \forall i \in [k]\}$ are disjoint convex bodies, by the separating hyperplane theorem (e.g., [BV04]), there exists nonzero $\mathbf{w} \in \mathbb{R}^k$ and $b \in \mathbb{R}$ such that for all $\mathbf{v} \in \mathbb{R}_+^k$, $\langle \mathbf{w}, \mathbf{v} \rangle \geq b$ and for all $\mathbf{y} \in H$, $\langle \mathbf{y}, \mathbf{w} \rangle \leq b$. Since \mathbf{w} is nonzero and \mathbb{R}_+^k is open, we must more strongly have that $\langle \mathbf{w}, \mathbf{v} \rangle > b$ and for all $\mathbf{v} \in \mathbb{R}_+^k$.

Taking the limit as $\mathbf{v} \rightarrow 0$ of $\langle \mathbf{y}, \mathbf{w} \rangle \leq b$ for all $\mathbf{y} \in H$, we have that $b \leq 0$. Further, since $(0, 0, \dots, 0) \in H$, we must have that $b = 0$ and since H is a vector space, $\langle \mathbf{y}, \mathbf{w} \rangle = 0$ for all $\mathbf{y} \in H$. Thus, \mathbf{w} is normal to H . Note that \mathbf{w} is nonzero and has all coordinates non-negative as $\langle \mathbf{v}, \mathbf{w} \rangle > 0$ for all \mathbf{v} in the positive orthant. \square

Proof of Lemma 4.9. Fix a pair of predicates $(P, Q) \in \Gamma$. It suffices to show that (P, Q) is ppp-definable from $\Gamma_{\text{AT}} \cup \Gamma_{\text{const}}$. If P has coordinates of fixed value, we can use a gadget reduction from Γ_{const} to simulate these values. Thus, we assume that P has non-trivial dependence in each coordinate, and thus we apply Lemma 4.12 to get that $Q \supseteq O_{\text{AT}}(P) = \{\text{sgn}(\mathbf{x} - \mathbf{y}) : \mathbf{x}, \mathbf{y} \in \text{Aff}(P), \forall i, x_i \neq y_i\}$. We may without loss of generality assume that $Q = \{\text{sgn}(\mathbf{x} - \mathbf{y}) : \mathbf{x}, \mathbf{y} \in \text{Aff}(P), \forall i, x_i \neq y_i\}$.

For every $\mathbf{x} \in \mathcal{B}^k \setminus Q$, we find \mathbf{w}, b such that

$$P_{\mathbf{w}, b} := \{\mathbf{v} \in \mathcal{B}^k : \langle \mathbf{w}, \mathbf{v} \rangle = b\}, \quad Q_{\mathbf{w}, b} := \mathcal{B}^k \setminus \{\text{sgn}(\mathbf{w}), -\text{sgn}(\mathbf{w})\}$$

satisfy $P \subseteq P_{\mathbf{w}, b}, Q \subseteq Q_{\mathbf{w}, b}$ and $\text{sgn}(\mathbf{w}) = \mathbf{x}$. By applying this for every $\mathbf{x} \in \mathcal{B}^k \setminus Q$, 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. $(P_i, Q_i) \in \Gamma_{\text{AT}}$ for every $i \in [L]$.
3. $\bigcap_{i \in [L]} Q_i = Q$.

This shows that (P, Q) is ppp-definable from $\{(P_1, Q_1), (P_2, Q_2), \dots, (P_L, Q_L)\} \subseteq \Gamma_{\text{AT}}$.

Henceforth, our goal is to show that for every $\mathbf{x} \in \mathcal{B}^k \setminus Q$, we can find \mathbf{w}, b such that $P_{\mathbf{w}, b} := \{x \in \mathcal{B}^k : \langle \mathbf{w}, x \rangle = b\}$ satisfies $P \subseteq P_{\mathbf{w}, b}$, $\langle \mathbf{w}, \text{sgn}(\mathbf{w}) \rangle > b$, $\langle \mathbf{w}, \text{sgn}(\mathbf{w}) \rangle > -b$ and $\text{sgn}(\mathbf{w}) = \mathbf{x}$. Without loss of generality, we can assume that $\mathbf{x} = (+1, +1, \dots, +1)$. Fix an arbitrary vector $\bar{\mathbf{x}} \in P$ such that $\bar{\mathbf{x}} \notin$

$\{(-1, -1, \dots, -1), (+1, +1, \dots, +1)\}$. Such a vector is guaranteed to exist as P does not contain \mathbf{x} and has non-trivial dependence on each coordinate. Let H be a subspace of \mathbb{R}^k defined as follows:

$$H := \{\mathbf{y} - \bar{\mathbf{x}} : \mathbf{y} \in \text{Aff}(P)\}$$

As $\mathbf{x} \notin O_{AT}(P)$, using Lemma 4.12, we get that for every $\mathbf{z} \in H$, $\text{sgn}(\mathbf{z}) \neq \mathbf{x}$, or in other words, there is no $\mathbf{z} \in H$ with $z_i > 0$ for all $i \in [k]$. Using Claim 4.13, we can obtain nonzero $\mathbf{w} \geq \mathbf{0}$ such that $\langle \mathbf{w}, \mathbf{y} \rangle = 0$ for all $\mathbf{y} \in H$. This shows that $\langle \mathbf{w}, \mathbf{y} \rangle = b$ for every $\mathbf{y} \in P$, where $b := \langle \mathbf{w}, \bar{\mathbf{x}} \rangle$ satisfies $\sum_i w_i > b$, $\sum_i w_i > -b$. \square

5 Unique Games based Hardness

In this section, we prove Theorem 1.4. First, we use the analysis of AT and MAJ polymorphisms for symmetric promise templates with folding and idempotence in [BG21] to show that we can relax Γ into one of five candidate PCSP types.

Lemma 5.1. *Let $\Gamma = (P, Q)$ be a Boolean promise template such that $\text{MAJ}_{L_1}, \text{AT}_{L_2} \notin \text{Pol}(\Gamma)$ for some odd integers L_1, L_2 . Then, there exists a Boolean template $\Gamma' = (P, Q)$ that is fipp-definable from Γ that is equal to one of the following:*

1. k is even, and $\Gamma_1 = (P, Q), P = \text{Ham}_k\{\frac{k}{2}\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \in \{1, k-1\}$.
2. k is odd, $\Gamma_2 = (P, Q), P = \text{Ham}_k\{l, \frac{k+1}{2}\}, Q = \text{Ham}_k\{0, 1, 2, \dots, k-1\}$, where $l \leq \frac{k-1}{2}$.
3. $\Gamma_3 = (P, Q), P = \text{Ham}_k\{l, k\}, Q = \text{Ham}_k\{1, 2, \dots, k\}$, where $l \neq 0, l \leq \frac{k-1}{2}$.
4. $\Gamma_4 = (P, Q), P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{0, k-1\}$ where $l \in \{1, 2, \dots, k-1\}, l \leq \frac{k-1}{2}$.
5. $\Gamma_5 = (P, Q), P = \text{Ham}_k\{1, k\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ for arbitrary $b \in \{0, 1, \dots, k\}$.

We defer the proof of Lemma 5.1 to Appendix B.

Recall that if a promise template Γ' is fipp-definable from another promise template Γ , if $\text{fPCSP}(\Gamma)$ has a polynomial time robust algorithm, then $\text{fPCSP}(\Gamma')$ has a polynomial time robust algorithm as well (Proposition 2.12). Thus, to show Theorem 1.4, it suffices to show Unique Games based hardness of obtaining robust algorithms for $\text{fPCSP}(\Gamma_i)$ for $i \in [5]$ from Lemma 5.1. We achieve these hardness results by showing integrality gaps for the basic SDP relaxation of each $\text{fPCSP}(\Gamma_i)$. 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 5.2 (Special case of [Rag08] for Boolean fPCSPs when the SDP is feasible). *Suppose that for a Boolean promise template Γ , there is a finite integrality gap for the basic SDP relaxation $\text{fPCSP}(\Gamma)$, i.e., there is a finite instance I of $\text{fPCSP}(\Gamma)$ on which the basic SDP relaxation is feasible but there is no assignment that weakly satisfies I . Then, there exists a constant $s < 1$ that is a function of Γ, I such that the following decision problem is NP-hard for sufficiently small $\varepsilon, \delta > 0$, assuming UGC. Given an instance Φ of $\text{fPCSP}(\Gamma)$, distinguish between the two cases:*

1. (Completeness.) *There exists an assignment that strongly satisfies $1 - \varepsilon$ fraction of the constraints in Φ .*

2. (Soundness.) No assignment weakly satisfies $s + \delta$ fraction of the constraints in Φ .

Remark 5.3. It is worth noting that Raghavendra does not work directly with the basic SDP relaxation of $\text{fiPCSP}(\Gamma)$ but rather with the basic SDP relaxation of $\text{PCSP}(\Gamma')$ for a suitable promise template Γ' . In particular, consider for each relation $(P, Q) \in \Gamma$ of arity k the variable set $V = \{x_1, \dots, x_k\}$. For every possible clause $C = (y_1, \dots, y_k) \in (V \cup \bar{V} \cup \mathcal{B})^k$ using variables $S \subseteq V$, we construct a corresponding relation pair (P_C, Q_C) of arity S such that

$$\begin{aligned} P_C &= \{z \in \mathcal{B}^S : C(z) \in P\} \\ Q_C &= \{z \in \mathcal{B}^S : C(z) \in Q\}. \end{aligned}$$

In particular, Γ' consists of every (P_C, Q_C) which can be defined from some $(P, Q) \in \Gamma$. The key observation is that every instance Ψ of $\text{fiPCSP}(\Gamma)$ can then be replaced by a suitable instance Ψ' of $\text{PCSP}(\Gamma')$ by replacing each use of (P, Q) with the corresponding (P_C, Q_C) . One can verify that each strong and weak assignment to Ψ has the same value as the corresponding assignment to Ψ' and that the SDP relaxations have the same value. Thus, Raghavendra's theorem does indeed apply.

Thus, to show Theorem 1.4, our goal is to show the existence of finite integrality gaps for the basic SDP relaxations of the Boolean promise templates in Lemma 5.1. To obtain such an integrality gap for the basic SDP relaxation of a PCSP, we study colorings of the n dimensional sphere \mathbb{S}^n that satisfy certain properties. We start by defining a few notations that we need.

Definition 5.4. Fix a predicate $P \subseteq \mathcal{B}^k$. We say that a tuple of vectors $V = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k), \mathbf{v}_i \in \mathbb{S}^n \forall i \in [k]$ are a P -configuration with respect to another vector $\mathbf{v}_0 \in \mathbb{S}^n$ if the tuple of vectors can be assigned to a set of literals in a constraint by the basic SDP relaxation of an instance of $\text{fiPCSP}(P, Q)$ with zero error, for some $Q \supseteq P$. In other words, there exists a probability distribution $\{\lambda(\mathbf{a}) : \mathbf{a} \in P\}$ supported on P that satisfies the following properties.

1. $0 \leq \lambda(\mathbf{a}) \leq 1$ for all $\mathbf{a} \in P$, and

$$\sum_{\mathbf{a} \in P} \lambda(\mathbf{a}) = 1.$$

2. First moments:

$$\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = \sum_{\mathbf{a} \in P} \lambda(\mathbf{a}) a_i \quad \forall i \in [k].$$

3. Second moments:

$$\langle \mathbf{v}_i, \mathbf{v}_{i'} \rangle = \sum_{\mathbf{a} \in P} \lambda(\mathbf{a}) a_i a_{i'} \quad \forall i, i' \in [k].$$

We now define the notion of a function respecting a Boolean fiPCSP . We refer to functions $f : \mathbb{S}^n \rightarrow \mathcal{B}$ as colorings of the sphere.

Definition 5.5. Fix a vector $\mathbf{v}_0 \in \mathbb{S}^n$. We say that a coloring of the sphere $f : \mathbb{S}^n \rightarrow \mathcal{B}$ is folded if $f(-\mathbf{v}) = -f(\mathbf{v})$ for every $\mathbf{v} \in \mathbb{S}^n$. We say that f is idempotent if $f(\mathbf{v}_0) = +1$ and $f(-\mathbf{v}_0) = -1$. We say that a coloring f respects the promise template (P, Q) with respect to a vector $\mathbf{v}_0 \in \mathbb{S}^n$ if the following condition holds. For every P -configuration $V = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k)$ with respect to \mathbf{v}_0 , we have that the colors of the vectors satisfy Q , i.e.,

$$(f(\mathbf{v}_1), f(\mathbf{v}_2), \dots, f(\mathbf{v}_k)) \in Q$$

More generally, we say that a coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$ respects a Boolean promise template Γ with respect to a vector \mathbf{v}_0 if it respects every predicate pair in $(P, Q) \in \Gamma$ with respect to \mathbf{v}_0 .

Our key observation is that the absence of such a folded sphere coloring respecting Γ for some finite n gives an integrality gap for the basic SDP relaxation of $\text{fiPCSP}(\Gamma)$.

Lemma 5.6. *For every Boolean promise template $\Gamma = \{(P_1, Q_1), \dots, (P_l, Q_l)\}$, the Basic SDP decides $\text{fiPCSP}(\Gamma)$ if and only if for every integer $n \geq 1$, there exists a folded, idempotent coloring $f^{(n)} : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ with respect to a vector $\mathbf{v}_0^{(n)}$.*

Proof. We slightly abuse the notation and say that a ‘‘partial’’ folded function $f : S \rightarrow \mathcal{B}, S \subseteq \mathbb{S}^n$ respects a Boolean promise template (P, Q) with respect to \mathbf{v}_0 if and only if for every P -configuration of vectors $V = (\mathbf{v}_1, \dots, \mathbf{v}_k)$ with respect to \mathbf{v}_0 with $\mathbf{v}_i \in S \forall i \in [k]$, $f(V) \in Q$. More generally, for a Boolean promise template Γ , $f : S \rightarrow \mathcal{B}$ respects Γ if and only if it respects every predicate pair in Γ . Via a compactness¹³ argument (e.g., like the De Bruijn-Erdos theorem [dBE51], for more details see Remark 7.13 of [BBKO21] or [CŽ23a]), we can infer that there is a folded coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$ respecting Γ with respect to \mathbf{v}_0 if and only if for every finite subset $S \subset \mathbb{S}^n$, there exists a folded coloring $f_S : S \rightarrow \mathcal{B}$ that respects Γ with respect to \mathbf{v}_0 .

First, assume that the Basic SDP decides $\text{fiPCSP}(\Gamma)$. Fix an arbitrary set of vectors $\mathbf{v}_0^{(n)} \in \mathbb{S}^n, n \in \mathbb{Z}^+$. Also pick a subset $T \subset \mathbb{S}^n$ such that for every $\mathbf{v} \in \mathbb{S}^n$, exactly one of \mathbf{v} or $-\mathbf{v}$ is in T . For any finite subset $S \subset \mathbb{S}^n$, we construct an instance Φ_S of Γ as follows. By increasing the size of S , we may assume that S is closed with respect to negation. The variable set of Φ_S is $\{x_{\mathbf{v}} : \mathbf{v} \in S \cap T\}$. We define the constraints as follows. For every $i \in [l]$ and every P_i -configuration $(\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k_i}) \in S^{k_i}$ with respect to $\mathbf{v}_0^{(n)}$, we add a constraint over (P_i, Q_i) using the literals $\ell_{\mathbf{v}_1}, \dots, \ell_{\mathbf{v}_{k_i}}$, where

$$\ell_{\mathbf{v}_j} = \begin{cases} +1 & \mathbf{v}_j = \mathbf{v}_0^{(n)} \\ -1 & \mathbf{v}_j = -\mathbf{v}_0^{(n)} \\ x_{\mathbf{v}_j} & \mathbf{v}_j \in T \\ \bar{x}_{-\mathbf{v}_j} & -\mathbf{v}_j \in T. \end{cases}$$

We have that $x_{\mathbf{v}} \mapsto \mathbf{v}$ is a basic SDP solution with zero error. Thus, there exists an assignment $f_{S \cap T} : S \cap T \rightarrow \mathcal{B}$ to the variables that weakly satisfies all the constraints in Φ_S , or equivalently, there exists folded, idempotent $f_S : S \rightarrow \mathcal{B}$ that respects the promise template Γ with respect to $\mathbf{v}_0^{(n)}$, where

$$f_S(\mathbf{v}) = \begin{cases} f_{S \cap T}(\mathbf{v}) & \mathbf{v} \in S \cap T \\ -f_{S \cap T}(\mathbf{v}) & \text{otherwise.} \end{cases}$$

Hence, there exists a folded, idempotent coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ with respect to $\mathbf{v}_0^{(n)}$ for every positive integer n .

Second, suppose that for every integer $n \geq 1$, there exists a folded, idempotent coloring $f^{(n)} : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ with respect to some vector $\mathbf{v}_0^{(n)}$. We seek to show that the Basic SDP decides $\text{fiPCSP}(\Gamma)$. Take an arbitrary instance Φ of $\text{fiPCSP}(\Gamma)$ such that there is a solution to Basic SDP with zero error. We solve the SDP relaxation of Φ and obtain a set of vectors \mathbf{v}_0 and $\mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{S}^n$ corresponding to the variables in Φ that satisfies all the constraints in Φ with zero error. By Proposition 2.17, there exists an orthogonal matrix $Q \in \mathbb{R}^{(n+1) \times (n+1)}$ for which $Q\mathbf{v}_0 = \mathbf{v}_0^{(n)}$. Since orthogonal matrices preserve the inner products of vectors, we have that multiplication on the left by Q is an automorphism of \mathbb{S}^n . Define $f^{(n)} : \mathbb{S}^n \rightarrow \mathcal{B}$ by $f^{(n)}(\mathbf{v}) = f^{(n)}(Q\mathbf{v})$. Since multiplication by an orthogonal matrix preserves inner products, we have that

¹³We assume the axiom of choice.

$f^{(n)}$ respects Γ with respect to \mathbf{v}_0 if and only if $f^{(n)}$ respects Γ with respect to $\mathbf{v}_0^{(n)}$. Therefore, the assignment $f^{(n)}(\mathbf{v}_i)$ to the variable u_i weakly satisfies all the constraints in Φ . Thus, for every instance Φ of $\text{fiPCSP}(\Gamma)$ with zero error on the basic SDP relaxation, there is an assignment that weakly satisfies all the constraints in Φ , or equivalently, the basic SDP decides $\text{fiPCSP}(\Gamma)$. \square

As a corollary, we get the following.

Corollary 5.7. *Let Γ be a Boolean promise template. Then, there is a finite integrality gap for the basic SDP relaxation of $\text{fiPCSP}(\Gamma)$ if and only if for some positive integer n , there exists no folded, idempotent coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ .*

Theorem 5.2 together with Corollary 5.7 shows that if for a Boolean promise template Γ does not admit a sphere coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ for some positive integer n , then, $\text{fiPCSP}(\Gamma)$ does not admit a polynomial time robust algorithm, assuming the Unique Games Conjecture. Thus, our goal is to show that the promise templates mentioned in Lemma 5.1 do not admit sphere coloring that respects them, and use Corollary 5.7 to prove Theorem 1.4.

In the rest of this section, we first prove a couple of lemmas regarding sphere Ramsey theory. Then, we show that the earlier mentioned promise templates Γ_{1-5} do not have folded sphere coloring respecting them using the sphere Ramsey results.

5.1 Sphere Ramsey theory

We start with some more notation. For a tuple of vectors $S = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k)$ with $\mathbf{v}_i \in \mathbb{S}^d$, we use $\rho(S)$ to denote the the smallest radius of a sphere that contains S as a subset.

$$\rho(S) := \min\{r : \exists \mathbf{c} \in \mathbb{R}^d, \|\mathbf{c} - \mathbf{v}_i\|_2 = r \forall i \in [k]\}.$$

Let $S_1 = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k)$, $S_2 = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k)$ with $\mathbf{u}_i \in \mathbb{S}^{d_1}$, $\mathbf{v}_i \in \mathbb{S}^{d_2}$ be two tuples with the same arity. We say that S_1 and S_2 are congruent if they have the same pairwise inner products, i.e., $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \langle \mathbf{v}_i, \mathbf{v}_j \rangle$ for all $i, j \in [k]$. Matoušek and Rödl [MR95] proved the following:

Theorem 5.8 ([MR95]). *Let $S = (\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k)$ be a tuple of affinely independent vectors with $\rho(S) < 1$. Then, for every positive integer $r \geq 2$, there exists an integer $n_0 := n_0(S, r)$ such that for every $n \geq n_0$, for every coloring $f : \mathbb{S}^n \rightarrow [r]$, there exists a tuple of vectors $S' = (\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k)$, $\mathbf{v}_i \in \mathbb{S}^n \forall i \in [k]$ that is congruent to S , and is monochromatic, i.e., $f(\mathbf{v}_i) = f(\mathbf{v}_j)$ for every $i, j \in [k]$.*

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

Lemma 5.9. *Fix an integer $k \geq 3$ and $r \geq 2$. There exists $n_0 := n_0(k, r)$ such that for every $n \geq n_0$ and coloring $f : \mathbb{S}^n \rightarrow [r]$ and $\gamma \in \mathbb{R}$ with $\frac{-1}{k-1} < \gamma < 1$, there exists a monochromatic set of vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} \subseteq \mathbb{S}^n$ such that $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \gamma$ for every $i \neq j$.*

Proof. Consider an arbitrary set $S = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\}$ of k unit vectors in \mathbb{S}^n such that $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \gamma$ for every $i \neq j$. Such a set S is guaranteed to exist when n is large enough. We show that the vectors are affinely independent: suppose for contradiction that there exists reals c_1, c_2, \dots, c_k not all zero, $\sum_i c_i = 0$ and $\sum_i c_i \mathbf{u}_i = 0$. We have

$$0 = \left\langle \mathbf{u}_1, \left(\sum_i c_i \mathbf{u}_i \right) \right\rangle = c_1 + \gamma(c_2 + \dots + c_k) = c_1 + \gamma(-c_1)$$

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

The set of vectors can be embedded on a sphere of radius strictly smaller than 1: let $\alpha \in \mathbb{R}$ such that $0 < \alpha < \frac{2}{k}$, and let $\mathbf{u}_s = \sum_{i \in [k]} \mathbf{u}_i$, $\mathbf{c} = \alpha \mathbf{u}_s$. We have

$$\|\mathbf{u}_s\|_2^2 = \sum_i \|\mathbf{u}_i\|_2^2 + 2 \sum_{i \neq j} \langle \mathbf{u}_i, \mathbf{u}_j \rangle = k + k(k-1)\gamma$$

Note that

$$\begin{aligned} \|\mathbf{u}_i - \mathbf{c}\|_2^2 &= \|\mathbf{u}_i\|_2^2 + \|\mathbf{c}\|_2^2 - 2\langle \mathbf{c}, \mathbf{u}_i \rangle \\ &= 1 + \alpha^2(k + k(k-1)\gamma) - 2\alpha(1 + (k-1)\gamma) \\ &= 1 - k(1 + (k-1)\gamma)\alpha \left(\alpha - \frac{2}{k} \right) \end{aligned}$$

which is strictly smaller than 1 when $0 < \alpha < \frac{2}{k}$. Thus, all the vectors are on a sphere centered at \mathbf{c} and radius strictly smaller than 1, implying that $\rho(S) < 1$. Now, we can use Theorem 5.8 on S and f to obtain the required set of vectors V . \square

While Theorem 5.8 is applicable to a wide range of sets S , we sometimes need to apply it to sets S that do not form a simplex or have $\rho(S) = 1$. Towards this, we use the ‘‘Spreads’’ based idea in [MR95] to obtain a version of Theorem 5.8 directly for certain sets S where Theorem 5.8 is not applicable.

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

$$\text{Spread-}s_n(\mathbf{a}, J) = \sum_{i=1}^k a_i \mathbf{e}_{j_i}$$

where $\mathbf{e}_1, \mathbf{e}_2, \dots, \mathbf{e}_n$ is an orthonormal basis of \mathbb{R}^n . For a set $I \subseteq [n]$, we let

$$\text{Spread}_n(\mathbf{a}, I) = \{\text{Spread-}s_n(\mathbf{a}, J) : J \subseteq I, |J| = k\}$$

Recall the hypergraph Ramsey theorem: for any n, k, r there exists N such that in any r -coloring of the hyperedges of the complete k -uniform hypergraph on N vertices, there exists an induced subgraph on n -vertices which is monochromatic. As a direct consequence, we get the following application to spreads.

Lemma 5.10. (*[MR95]*) *For every $\mathbf{a} \in \mathbb{R}^k, n, k$, there exists N such that in any coloring $f : \text{Spread}_N(\mathbf{a}, [N]) \rightarrow [r]$, there exists I with $|I| = n$ such that $\text{Spread}_N(\mathbf{a}, I)$ is monochromatic with respect to f , i.e., $\exists p \in [r]$ such that $f(\mathbf{v}) = p$ for all $\mathbf{v} \in \text{Spread}_N(\mathbf{a}, I)$.*

Lemma 5.10 implies the following immediately.

Corollary 5.11. *Let $U = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\}$ be a set of k unit vectors such that $\mathbf{u}_i \in \text{Spread}_N(\mathbf{a}, [N])$ for all $i \in [k]$ for an integer N , and a vector $\mathbf{a} \in \mathbb{R}^N$ with $\|\mathbf{a}\|_2 = 1$. Then there exists $n_0 := n_0(U, \mathbf{a}, N)$ such that for every $n \geq n_0, r$, for every sphere coloring $f : \mathbb{S}^n \rightarrow [r]$, there exists a set of k vectors $V = \{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ that are all colored the same, and $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \langle \mathbf{u}_i, \mathbf{u}_j \rangle$ for every $i, j \in [k]$.*

We use Corollary 5.11 to prove a lemma regarding sphere colorings. For ease of notation, we call a set of k unit vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ to be k -regular if $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = -\frac{1}{k-1}$ for every $i \neq j$.

Lemma 5.12. Fix an integer $k \geq 2$. There exists $n_0 := n_0(k)$ such that for every $n \geq n_0$ and folded coloring $f: \mathbb{S}^n \rightarrow \mathcal{B}$, there exist a k -regular set of vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} \subseteq \mathbb{S}^n$ such that exactly $k - 1$ vectors in V are colored -1 .

Proof. We construct a set of k unit vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ in $\text{Spread}_N(\mathbf{a}, [N])$ such that the set of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k-1}, -\mathbf{v}_k\}$ is a k -regular set, where N and \mathbf{a} depend only on k , and $\|\mathbf{a}\|_2 = 1$. Using Corollary 5.11, we can infer that in the coloring f , there exist k vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ that are all assigned the same color, such that $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k-1}, -\mathbf{v}_k\}$ is a k -regular set. As f is folded, this implies that there is a k -regular set in which exactly $k - 1$ vectors are assigned the color -1 .

Thus our goal is to construct k unit vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ in $\text{Spread}_N(\mathbf{a}, [N])$ such that the set of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k-1}, -\mathbf{v}_k\}$ is a k -regular set. Or equivalently, we construct the vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{k-1}$ in $\text{Spread}_N(\mathbf{a}, [N])$ and \mathbf{v}_k in $\text{Spread}_N(-\mathbf{a}, [N])$ such that $\{\mathbf{v}_1, \dots, \mathbf{v}_k\}$ is a k -regular set. We set $\gamma = \frac{1}{\sqrt{2(k-1)}}$ and $\mathbf{a} = (\gamma, -\gamma, \gamma, -\gamma, \dots, -\gamma) \in \mathbb{R}^{2(k-1)}$. We set $\mathbf{v}_i = \text{Spread}_n(\mathbf{a}, J_i), i \in [k-1], \mathbf{v}_k = \text{Spread}(-\mathbf{a}, J_k)$ 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 on k . 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:

$$\begin{aligned}\mathbf{v}_1 &= (\gamma, -\gamma) \\ \mathbf{v}_2 &= (-\gamma, \gamma)\end{aligned}$$

where $\gamma = \frac{1}{\sqrt{2}}$. Note that the above two vectors are a 2-regular set, and letting $\mathbf{a} = (\gamma, -\gamma)$, we have $\mathbf{v}_1 \in \text{Spread}_2(\mathbf{a}, [2])$, and $\mathbf{v}_2 \in \text{Spread}_2(-\mathbf{a}, [2])$. Now, suppose that J_1, J_2, \dots, J_k, N are such that $\mathbf{v}_i = \text{Spread}_N(\mathbf{a}, J_i), i \in [k-1], \mathbf{v}_k = \text{Spread}_N(-\mathbf{a}, J_k)$ satisfy the property that $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ is a k -regular set with $\mathbf{a} = (\gamma, -\gamma, \dots, \gamma, -\gamma) \in \mathbb{R}^{2(k-1)}, \gamma = \frac{1}{\sqrt{2(k-1)}}$. We construct J'_1, \dots, J'_{k+1} such that $\mathbf{v}'_i = \text{Spread}_{N'}(\mathbf{a}', J'_i)$ for all $i \in [k], \mathbf{v}'_{k+1} = \text{Spread}_{N'}(-\mathbf{a}', J'_{k+1})$ satisfy the property that $\{\mathbf{v}'_1, \mathbf{v}'_2, \dots, \mathbf{v}'_{k+1}\}$ is a $(k+1)$ -regular set with $\mathbf{a}' = (\gamma', -\gamma', \dots, \gamma', -\gamma') \in \mathbb{R}^{2k}, \gamma' = \frac{1}{\sqrt{2k}}$.

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

$$J'_i = J_i \cup \{N + 2i, N + 2i + 1\}$$

This ensures that $\langle \mathbf{v}'_i, \mathbf{v}'_j \rangle = -(\gamma')^2$ for every $i, j \in [k-1], i \neq j$.

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

$$J_{k+1} = J_k \cup \{N + 1, N + 2k\}$$

This ensures that $\langle \mathbf{v}'_i, \mathbf{v}'_{k+1} \rangle = -(\gamma')^2$ for every $i \in [k-1]$.

3. Finally, we set J_k .

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

This ensures that $\langle \mathbf{v}'_i, \mathbf{v}'_k \rangle = -(\gamma')^2$ for every $i \in [k+1], i \neq k$.

We illustrate our construction by obtaining the vectors for the case when $k = 3$ and $k = 4$:

$$\begin{aligned}\mathbf{v}_1 &= (\alpha, -\alpha, 0, \alpha, -\alpha, 0) \\ \mathbf{v}_2 &= (0, 0, \alpha, -\alpha, \alpha, -\alpha) \\ \mathbf{v}_3 &= (-\alpha, \alpha, -\alpha, 0, 0, \alpha)\end{aligned}$$

$$\begin{aligned}
\mathbf{v}_1 &= (\beta, -\beta, 0, \beta, -\beta, 0, 0, \beta, -\beta, 0, 0, 0) \\
\mathbf{v}_2 &= (0, 0, \beta, -\beta, \beta, -\beta, 0, 0, 0, \beta, -\beta, 0) \\
\mathbf{v}_3 &= (0, 0, 0, 0, 0, 0, \beta, -\beta, \beta, -\beta, \beta, -\beta) \\
\mathbf{v}_4 &= (-\beta, \beta, -\beta, 0, 0, \beta, -\beta, 0, 0, 0, 0, \beta)
\end{aligned}$$

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

As the pairwise inner product of every pair in $\{\mathbf{v}'_1, \mathbf{v}'_2, \dots, \mathbf{v}'_{k+1}\}$ is equal to $-(\gamma')^2 = -\frac{1}{k}$, we get that these set of vectors are a $(k+1)$ -regular set, completing the inductive proof. \square

5.2 Absence of sphere coloring via sphere Ramsey theory

First, we show the absence of sphere coloring respecting Γ_1 using Lemma 5.12.

Lemma 5.13. *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 $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects $\Gamma_1 = (P, Q)$, $P = \text{Ham}_k\{\frac{k}{2}\}$, $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \in \{1, k-1\}$.*

Proof. Consider a large integer n and suppose for the sake of contradiction that there is a folded function $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ_1 with respect to a vector $\mathbf{v}_0 \in \mathbb{S}^n$. We get the P -configuration of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k$ where we set $\lambda(\mathbf{a}) = \frac{1}{|P|}$ for every $\mathbf{a} \in P$ in Definition 5.4. The vectors satisfy the following properties.

1. (First moments.) $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = 0$ for every $i \in [k]$.
2. (Second moments.) $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \frac{2\binom{k/2}{2} - \frac{k^2}{4}}{\binom{k}{2}} = \frac{-1}{k-1}$ for every $i \neq j \in [k]$.

Our goal is to show that there is a P -configuration of such vectors such that exactly b of them are colored $+1$ according to f . Consider the set of vectors

$$\mathbf{v}_0^\perp := \{\mathbf{u} \in \mathbb{S}^n : \langle \mathbf{u}, \mathbf{v}_0 \rangle = 0\}$$

Using Lemma 5.12, we can obtain a set of k vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k \in \mathbf{v}_0^\perp$ such that $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \frac{-1}{k-1}$ and exactly $k-1$ of $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\}$ are colored -1 .

$$|\{i : f(\mathbf{u}_i) = -1\}| = k-1.$$

Note that both $\{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\}$ and $\{-\mathbf{u}_1, -\mathbf{u}_2, \dots, -\mathbf{u}_k\}$ are P -configurations with respect to \mathbf{v}_0 . Since f is folded, in at least one of these two P -configurations, there are exactly b vectors that are colored $+1$, a contradiction. \square

We show the absence of sphere coloring respecting Γ_2, Γ_3 , and Γ_4 using Lemma 5.9.

Lemma 5.14. *Fix an odd integer $k \geq 3$ and integer $l : 0 \leq l \leq \frac{k-1}{2}$. There exists an integer n_0 such that for every $n \geq n_0$, there is no folded $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects $\Gamma_2 = (P, Q)$, $P = \text{Ham}_k\{l, \frac{k+1}{2}\}$, $Q = \text{Ham}_k\{0, 1, 2, \dots, k-1\}$.*

Proof. Consider a large integer n and suppose for contradiction that there is a folded function $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ_2 with respect to a vector $\mathbf{v}_0 \in \mathbb{S}^n$. The P -configuration that we consider is a set of vectors

$\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k$ that are obtained by setting $\lambda(\mathbf{a})$ in Definition 5.4 as follows. We first sample an integer $t \in \{l, \frac{k+1}{2}\}$ as below.

$$t = \begin{cases} l, & \text{with probability } \frac{1}{1-s}. \\ \frac{k+1}{2}, & \text{with probability } \frac{-s}{1-s}. \end{cases}$$

where $s = l - (k - l) < 0$. The probability distribution λ is obtained by sampling a uniform element of $\text{Ham}_k\{t\}$. In other words, we have

$$\lambda(\mathbf{a}) := \begin{cases} \frac{1}{(1-s) \binom{k}{l}}, & \text{if } \mathbf{a} \in \text{Ham}_k\{l\}. \\ \frac{-s}{(1-s) \binom{k}{\frac{k+1}{2}}}, & \text{else if } \mathbf{a} \in \text{Ham}_k\{\frac{k+1}{2}\}. \\ 0, & \text{otherwise.} \end{cases}$$

We obtain the following properties:

1. (First moments). By symmetry, $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = \langle \mathbf{v}_j, \mathbf{v}_0 \rangle$ for all $i, j \in [k]$.

We have $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = \frac{1}{k} \langle (\sum_{i \in [k]} \mathbf{v}_i), \mathbf{v}_0 \rangle = \frac{1}{k} \left(\frac{1}{1-s} (l - (k-l)) + \frac{-s}{1-s} 1 \right) = 0$ for every $i \in [k]$.

2. (Second moments). By symmetry of variables, we get that $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \gamma$ for every $i \neq j$, for some $\gamma := \gamma(k, l)$. We have

$$k + k(k-1)\gamma = \left\| \sum_i \mathbf{v}_i \right\|^2 = \sum_{\mathbf{a} \in P} \lambda(\mathbf{a}) \left(\sum_{i \in [k]} a_i \right)^2 > 0.$$

Thus, we get that $\frac{-1}{k-1} < \gamma < 1$.

Now, restricting ourselves to the vectors in \mathbb{S}^n that are orthogonal to \mathbf{v}_0 , and using Lemma 5.9, we get that there exists a P -configuration of vectors that are all colored the same. By taking the negation of these vectors if needed, we get our required claim. \square

Lemma 5.15. 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 $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects $\Gamma_3 = (P, Q)$, $P = \text{Ham}_k\{l, k\}$, $Q = \text{Ham}_k\{1, 2, \dots, k\}$, where $l \neq 0, l \leq \frac{k-1}{2}$.

Proof. Consider a large integer n and suppose for contradiction that there is a folded function $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ_3 with respect to a vector $\mathbf{v}_0 \in \mathbb{S}^n$. We pick the P -configuration along the same lines as in Lemma 5.14. We sample $t \in \{l, k\}$ with

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

where $s = l - (k - l) < 0$. As before, the probability distribution λ is obtained by sampling a uniform element of $\text{Ham}_k\{t\}$. We get

1. (First moments). $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = \frac{1}{k} \left\{ \left(\frac{k}{k-s} \right) \frac{s}{k} + \left(\frac{-s}{k-s} \right) 1 \right\} = 0$ for every $i \in [k]$.

2. (Second moments). As in Lemma 5.14, we have $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \gamma$ for every $i \neq j$, for some $\gamma := \gamma(k, l)$ with

$$k + k(k-1)\gamma = \left\| \sum_i \mathbf{v}_i \right\|^2 = \sum_{\mathbf{a} \in P} \lambda(\mathbf{a}) \left(\sum_{i \in [k]} a_i \right)^2 > 0.$$

Thus, we get that $\frac{-1}{k-1} < \gamma < 1$.

We restrict ourselves to vectors in \mathbb{S}^n that are orthogonal to \mathbf{v}_0 , and applying Lemma 5.9, we get that for any coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$, there is a monochromatic P -configuration that we described. By negating the vectors if needed, we get our required proof. \square

For the proof of the next case Γ_4 we will need Lemma 5.9 applied to 4-colorings of the sphere.

Lemma 5.16. *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 $f : \mathbb{S}^n \rightarrow \{0, 1\}$ that is folded and respects the promise template $\Gamma_4 = (P, Q), P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{0, k-1\}$.*

Proof. As before, consider a large integer n and suppose for contradiction that there is a folded function $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects Γ_4 with respect to a vector $\mathbf{v}_0 \in \mathbb{S}^n$. We partition the predicate P into P_1 and P_{-1} depending on the value of the first element, i.e.,

$$P_i = \{\mathbf{x} \in P : x_1 = i\}, i \in \mathcal{B}$$

We pick the probability distribution λ as follows: sample i from \mathcal{B} uniformly at random, then, sample a uniformly random element from P_i . We get

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

$$\langle \mathbf{v}_1, \mathbf{v}_0 \rangle = 0$$

By using the symmetry of the rest of the variables and the fact that $\sum_{i=1}^k \langle \mathbf{v}_i, \mathbf{v}_0 \rangle = l - (k-l) = 2l - k$, we get that

$$\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = \frac{2l-k}{k-1} \forall i \in \{2, 3, \dots, k\}.$$

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

2. (Second moments). We have

$$\begin{aligned} \left\langle \mathbf{v}_1, \left(\sum_{i=2}^k \mathbf{v}_i \right) \right\rangle &= \frac{1}{2} (1 \cdot (2l-k-1)) + \frac{1}{2} ((-1) \cdot (2l-k+1)) \\ &= -1 \end{aligned}$$

Thus, we get

$$\langle \mathbf{v}_1, \mathbf{v}_i \rangle = \frac{-1}{k-1} \forall i \in \{2, 3, \dots, k\}.$$

Note that we have

$$\left\| \sum_{i=1}^k \mathbf{v}_i \right\|_2^2 = (2l-k)^2.$$

Using this, we get that

$$\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \frac{(2l-k)^2 - (k-2)}{(k-1)(k-2)} \forall i, j \in \{2, 3, \dots, k\}, i \neq j$$

For ease of notation, let $\beta = \frac{(2l-k)^2 - (k-2)}{(k-1)(k-2)}$.

Our goal is to show that there exists n_0 such that for every $n \geq n_0$, for every folded sphere coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$ and $\mathbf{v}_0 \in \mathbb{S}^n$, there exists a set of k vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\}$ that satisfy the above first and second moments, and exactly b vectors in V are colored $+1$, where $b \in \{0, k-1\}$.

Let $\gamma = \frac{\beta - \alpha^2}{1 - \alpha^2}$. We expand this as follows:

$$\begin{aligned} \gamma &= \frac{\frac{(2l-k)^2 - (k-2)}{(k-1)(k-2)} - \frac{(2l-k)^2}{(k-1)^2}}{1 - \alpha^2} \\ &= \frac{\frac{(2l-k)^2}{(k-1)^2(k-2)} - \frac{1}{k-1}}{1 - \alpha^2} \\ &= \frac{\frac{\alpha^2}{k-2} - \frac{1}{k-1}}{1 - \alpha^2} = \frac{\frac{1 - (1 - \alpha^2)}{k-2} - \frac{1}{k-1}}{1 - \alpha^2}. \end{aligned}$$

Simplifying further, we obtain:

$$\gamma = \frac{1}{(1 - \alpha^2)(k-1)(k-2)} - \frac{1}{k-2} > -\frac{1}{k-2}. \quad (8)$$

We apply Lemma 5.9 on the following coloring of the sphere. For a vector $\mathbf{u} \in \mathbb{S}^n$ such that $\langle \mathbf{u}, \mathbf{v}_0 \rangle = 0$, let $f' : \mathbb{S}^{n-1} \rightarrow \mathcal{B}^2$ be defined as

$$f'(\mathbf{u}) = \left(f\left(\alpha \mathbf{v}_0 + \sqrt{1 - \alpha^2} \mathbf{u}\right), f\left(\alpha \mathbf{v}_0 - \sqrt{1 - \alpha^2} \mathbf{u}\right) \right)$$

Using Lemma 5.9 on f' combined with the fact that $\gamma > -\frac{1}{k-2}$ obtained from Equation (8), we can infer that there exist $k-1$ unit vectors $\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_{k-1} \in \mathbb{S}^{n-1}$ such that $\langle \mathbf{u}_i, \mathbf{v}_0 \rangle = 0$ for all i , $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \gamma$ for all $i \neq j$ and $f'(\mathbf{u}_i) = f'(\mathbf{u}_j)$ for all $i \neq j, i, j \in [k-1]$.

We define $\mathbf{v}_1^{(1)}, \mathbf{v}_2^{(1)}, \dots, \mathbf{v}_k^{(1)}, \mathbf{v}_1^{(2)}, \dots, \mathbf{v}_k^{(2)}$ as follows. For $i \in \{2, 3, \dots, k\}$, we let

$$\begin{aligned} \mathbf{v}_i^{(1)} &= \alpha \mathbf{v}_0 + \sqrt{1 - \alpha^2} \mathbf{u}_{i-1} \\ \mathbf{v}_i^{(2)} &= \alpha \mathbf{v}_0 - \sqrt{1 - \alpha^2} \mathbf{u}_{i-1} \end{aligned}$$

We let

$$\mathbf{v}_1^{(1)} = -\frac{\sum_{i=1}^{k-1} \mathbf{u}_i}{\left\| \sum_{i=1}^{k-1} \mathbf{u}_i \right\|}$$

and $\mathbf{v}_1^{(2)} = -\mathbf{v}_1^{(1)}$. We now prove that the set of vectors $\mathbf{v}_1^{(1)}, \dots, \mathbf{v}_k^{(1)}$ and the set of vectors $\mathbf{v}_1^{(2)}, \dots, \mathbf{v}_k^{(2)}$ are a P -configuration with first and second moments as computed earlier, where we sampled i from \mathcal{B} uniformly at random and set the probability distribution λ as the uniform distribution over P_i .

1. (First moments). As $\langle \mathbf{u}_i, \mathbf{v}_0 \rangle = 0$ for all $i \in [k-1]$, we get that

$$\langle \mathbf{v}_1^{(1)}, \mathbf{v}_0 \rangle = \langle \mathbf{v}_1^{(2)}, \mathbf{v}_0 \rangle = 0$$

and

$$\langle \mathbf{v}_i^{(1)}, \mathbf{v}_0 \rangle = \langle \mathbf{v}_i^{(2)}, \mathbf{v}_0 \rangle = \alpha \quad \forall i \in \{2, \dots, k\}$$

2. (Second moments). We have $\forall i \in \{2, 3, \dots, k\}$,

$$\begin{aligned} \langle \mathbf{v}_1^{(1)}, \mathbf{v}_i^{(1)} \rangle &= -\frac{\langle (\sum_{j=1}^{k-1} \mathbf{u}_j), (\alpha \mathbf{v}_0 + \sqrt{1-\alpha^2} \mathbf{u}_{i-1}) \rangle}{\left\| \sum_{j=1}^{k-1} \mathbf{u}_j \right\|} \\ &= \frac{\sqrt{1-\alpha^2}}{k-1} \frac{\langle (\sum_{j=1}^{k-1} \mathbf{u}_j), (\sum_{j=1}^{k-1} \mathbf{u}_j) \rangle}{\left\| \sum_{j=1}^{k-1} \mathbf{u}_j \right\|} \\ &= -\frac{\sqrt{1-\alpha^2}}{k-1} \left\| \sum_{j=1}^{k-1} \mathbf{u}_j \right\| \\ &= -\frac{\sqrt{1-\alpha^2}}{k-1} \sqrt{k-1 + 2 \frac{(k-1)(k-2)}{2} \gamma} \\ &= -\frac{\sqrt{1-\alpha^2}}{k-1} \sqrt{\frac{1}{1-\alpha^2}} \text{ Using Equation (8)} \\ &= \frac{-1}{k-1}. \end{aligned}$$

Furthermore, $\forall i, j \in \{2, 3, \dots, k\}, i \neq j$,

$$\begin{aligned} \langle \mathbf{v}_i^{(1)}, \mathbf{v}_j^{(1)} \rangle &= \langle (\alpha \mathbf{v}_0 + \sqrt{1-\alpha^2} \mathbf{v}_i), (\alpha \mathbf{v}_0 + \sqrt{1-\alpha^2} \mathbf{v}_j) \rangle \\ &= \alpha^2 + (1-\alpha^2) \gamma \\ &= \beta. \end{aligned}$$

Similarly, we have

$$\begin{aligned} \langle \mathbf{v}_1^{(2)}, \mathbf{v}_i^{(2)} \rangle &= \frac{\langle (\sum_{j=1}^{k-1} \mathbf{u}_j), (\alpha \mathbf{v}_0 - \sqrt{1-\alpha^2} \mathbf{u}_{i-1}) \rangle}{\left\| \sum_{j=1}^{k-1} \mathbf{u}_j \right\|} = \frac{-1}{k-1} \quad \forall i \in \{2, 3, \dots, k\}. \\ \langle \mathbf{v}_i^{(2)}, \mathbf{v}_j^{(2)} \rangle &= \langle (\alpha \mathbf{v}_0 - \sqrt{1-\alpha^2} \mathbf{v}_i), (\alpha \mathbf{v}_0 - \sqrt{1-\alpha^2} \mathbf{v}_j) \rangle \\ &= \alpha^2 + (1-\alpha^2) \gamma = \alpha^2 + (1-\alpha^2) \gamma = \beta \quad \forall i, j \in \{2, 3, \dots, k\}, i \neq j \end{aligned}$$

Thus, both the set of vectors $\mathbf{v}_1^{(1)}, \mathbf{v}_2^{(1)}, \dots, \mathbf{v}_k^{(1)}$ and the set of vectors $\mathbf{v}_1^{(2)}, \mathbf{v}_2^{(2)}, \dots, \mathbf{v}_k^{(2)}$ are P -configurations with respect to \mathbf{v}_0 . As $f'(\mathbf{u}_1) = f'(\mathbf{u}_2) = \dots = f'(\mathbf{u}_{k-1})$, we can infer that $f(\mathbf{v}_2^{(1)}) = f(\mathbf{v}_3^{(1)}) = \dots = f(\mathbf{v}_k^{(1)})$ and $f(\mathbf{v}_2^{(2)}) = f(\mathbf{v}_3^{(2)}) = \dots = f(\mathbf{v}_k^{(2)})$. Furthermore, as $\mathbf{v}_1^{(1)} = -\mathbf{v}_1^{(2)}$ and f is folded, we can infer that $f(\mathbf{v}_1^{(1)}) = -f(\mathbf{v}_1^{(2)})$. Thus, there exists $p \in \{1, 2\}$ such that $f(\mathbf{v}_1^{(p)}) = -1$. Thus, there are either 0 or $k-1$ vectors among $\mathbf{v}_1^{(p)}, \mathbf{v}_2^{(p)}, \dots, \mathbf{v}_k^{(p)}$ that are colored $+1$ according to f , contradicting the fact that f respects the PCSP (P, Q) . \square

5.3 Absence of sphere coloring via connectivity of configurations

Finally, we show the absence of sphere coloring for Γ_5 using a connectivity lemma.

Lemma 5.17. *Fix integers $k \geq 3, b \in \{0, 1, \dots, k\} \setminus \{1, k\}$. There exists an integer n_0 such that for every $n \geq n_0$, there does not exist coloring $f : \mathbb{S}^n \rightarrow \{0, 1\}$ that is folded and respects the promise template $\Gamma_5 = (P, Q)$, $P = \text{Ham}_k\{1, k\}$, $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$.*

We dedicate the rest of the section to proving Lemma 5.17. We pick the configuration of vectors along the same lines as in Lemma 5.14. Fix $\mathbf{v}_0 \in \mathbb{S}^n$. The P -configuration that we study is a set of vectors $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k$ that is obtained by first sampling $t \in \{1, k\}$ such that

$$t = \begin{cases} 1 & \text{with probability } \frac{k}{2k-2} \\ k & \text{with probability } \frac{k-2}{2k-2} \end{cases}$$

Then, we sample a uniform element from $\text{Ham}_k\{t\}$. We get the following properties:

1. (First moments). $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = \left(\frac{k}{2k-2}\right) \frac{2-k}{k} + \left(\frac{k-2}{2k-2}\right) 1 = 0$ for every $i \in [k]$.
2. (Second moments). For every $i \neq j \in [k]$, we get

$$\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \left(\frac{k}{2k-2}\right) \frac{\binom{k-1}{2} - (k-1)}{\binom{k}{2}} + \left(\frac{k-2}{2k-2}\right) 1 = \frac{k-3}{k-1}$$

For ease of notation, let $\alpha = \frac{k-3}{k-1}$. We restrict our attention to vectors in the unit sphere \mathbb{S}^n that are orthogonal to \mathbf{v}_0 . This allows us to focus on P -configurations which are sets of k unit vectors with pairwise inner products equal to α . More formally, we define such a set of vectors as an α -configuration.

Definition 5.18. *A set of vectors $V = \{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k\} \subset \mathbb{S}^n$ is called an α -configuration if the vectors \mathbf{v}_i satisfy:*

$$\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \alpha \quad \text{for all } i \neq j.$$

Given a folded sphere coloring f , our objective is to prove that there exists an α -configuration V such that exactly b of the vectors in V are assigned the value $+1$.

Unlike the earlier studied promise template, here, the setting when $b = 0$ is relatively straightforward, simply because $\alpha \geq 0$. $\alpha \geq 0$ 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 α -configurations that are all colored $+1$, and also α -configurations 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 k of them are colored $+1$ according to f . Thus, we can pick k pairwise orthogonal vectors from this set where exactly b of them are colored $+1$ according to f . Henceforth, we assume that $\alpha > 0$.

To show that there are α -configurations that have exactly b vectors that are colored $+1$, we show a connectivity lemma (Lemma 5.21) where we prove that between any two α -configurations, there exists a

path using $O_{k,\alpha}(1)$ α -configurations where we change a single vector at each step in the path. As there is an α -configuration where all k vectors are colored $+1$, and the α -configuration obtained by negating these vectors where all the vectors are colored -1 , the connectivity lemma then shows that for every $b \in \{0, 1, \dots, k\}$, there exists an α -configuration that has exactly b vectors that are colored $+1$.

We first prove a simplified version of the connectivity lemma that we use to prove Lemma 5.21.

Lemma 5.19. *Given an α -configuration $U = \{\mathbf{u}_1, \mathbf{u}_2, \dots, \mathbf{u}_k\} \subseteq \mathbb{S}^n$, and a unit vector $\mathbf{w} \in \mathbb{S}^n$ that is orthogonal to each vector in U , there exists $L := L(k, \alpha)$ and a set of α -configurations V_1, V_2, \dots, V_L such that*

1. *The consecutive configurations differ in a single vector i.e., $|V_i \cap V_{i+1}| = k - 1$ for every $i \in [L - 1]$.*
2. *Final configuration contains \mathbf{w} , i.e., $\mathbf{w} \in V_L$, and the initial configuration V_1 is equal to U .*

Proof. We prove the lemma by studying the inner product of \mathbf{w} with an α -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 \mathbf{w} eventually reaches all α s. Towards this end, for an α -configuration V , we define $I(V, \mathbf{w})$ as the Gram matrix of $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_k, \mathbf{w}\}$.

$$I(V, \mathbf{w})_{i,j} = \begin{cases} \langle \mathbf{v}_i, \mathbf{v}_j \rangle & \text{if } 1 \leq i, j \leq k. \\ \langle \mathbf{v}_i, \mathbf{w} \rangle & \text{if } i = k + 1, 1 \leq j \leq k. \\ \langle \mathbf{w}, \mathbf{v}_j \rangle & \text{if } j = k + 1, 1 \leq i \leq k. \\ \langle \mathbf{w}, \mathbf{w} \rangle = 1 & \text{if } i = j = k + 1. \end{cases}$$

Starting with $I(V, \mathbf{w})$ where $V = U$, our goal is to change one vector in V at a time so that we eventually reach a configuration where the last column in $I(V, \mathbf{w})$ is equal to $(\alpha, \alpha, \dots, \alpha, 1)$. Note that changing one vector in V corresponds to changing a single value in the last column (and the corresponding value in the last row) in $I(V, \mathbf{w})$. 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 $I(V, \mathbf{w})$, we obtain a new matrix that is $I(V', \mathbf{w})$ with V' being different from V only in a single vector, as long as the new matrix is positive semidefinite.

Claim 5.20. *Suppose that A is a $m \times m$ real symmetric positive semidefinite matrix with $A = U^T U$ with $U = [\mathbf{u}_1 \mathbf{u}_2 \dots \mathbf{u}_m]$ where $\mathbf{u}_i \in \mathbb{R}^n$ with $n \geq m$, and A' is another real symmetric positive semidefinite matrix such that A' and A differ only in the entries $(1, 2)$ and $(2, 1)$. Then, there exists $\mathbf{u}'_1 \in \mathbb{R}^n$ such that for $U' := [\mathbf{u}'_1 \mathbf{u}_2 \dots \mathbf{u}_m]$ we have that $A' = (U')^T U'$.*

Proof. As A' is a positive semidefinite matrix, there exist $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_m \in \mathbb{R}^n$ such that $A' = V^T V$ where $V = [\mathbf{v}_1 \mathbf{v}_2 \dots \mathbf{v}_m]$. Observe that $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \langle \mathbf{v}_i, \mathbf{v}_j \rangle$ for all $i, j \in \{2, \dots, m\}$. Therefore, by Proposition 2.17, there exists an orthogonal matrix $Q \in \mathbb{R}^{n \times n}$ such that $\mathbf{v}_i = Q\mathbf{u}_i$ for all $i \in \{2, \dots, m\}$. Thus, $A' = (VQ^T)^T (Q^T V)$, where the last $m - 1$ columns of $Q^T V$ are $\mathbf{u}_2, \dots, \mathbf{u}_m$. We set \mathbf{u}'_1 to be the first column of $Q^T V$. \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. $M_1 = I(U, \mathbf{w})$.
2. The diagonal entries of M_i for every $i \in [L]$ are all equal to 1.

3. All the off-diagonal entries of M_L are equal to α .
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 $\varepsilon \geq 0$, $0 \leq \gamma \leq \alpha$, and $d \in [k]$, we define the $(k+1) \times (k+1)$ matrix $M(\gamma, \varepsilon, d)$ as follows:

$$M(\gamma, \varepsilon, d)_{i,j} = \begin{cases} 1, & \text{if } i = j. \\ \alpha, & \text{if } 1 \leq i, j \leq k. \\ \gamma + \varepsilon, & \text{if } i = k+1, 1 \leq j \leq d \text{ or } j = k+1, 1 \leq i \leq d. \\ \gamma, & \text{if } i = k+1, d+1 \leq j \leq k \text{ or } j = k+1, d+1 \leq i \leq k. \end{cases}$$

Note that $I_{U,w} = M(0, 0, k)$, and our goal M_L is equal to $M(\alpha, 0, k)$. We consider the below sequence of positive semidefinite matrices $M(0, 0, k), M(0, \varepsilon, 1), M(0, \varepsilon, 2), \dots, M(0, \varepsilon, k), M(\varepsilon, \varepsilon, 1), M(\varepsilon, \varepsilon, 2), \dots, M(\varepsilon, \varepsilon, k), M(2\varepsilon, \varepsilon, 1), \dots, M(\alpha - \varepsilon, \varepsilon, k)$. 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 $\varepsilon \leq \frac{1-\alpha}{k}$, $M(\gamma, \varepsilon, d)$ is positive semidefinite for every $d \in [k], 0 \leq \gamma \leq \alpha$. This follows from a simple calculation.

$$\begin{aligned} x^T M(\gamma, \varepsilon, d)x &= \gamma \left(\sum_{i=1}^{k+1} x_i \right)^2 + (\alpha - \gamma) \left(\sum_{i=1}^k x_i \right)^2 + (1 - \alpha) \sum_{i=1}^k x_i^2 + (1 - \gamma)x_{k+1}^2 + 2\varepsilon(x_{k+1}) \left(\sum_{i=1}^d x_i \right) \\ &\geq (1 - \alpha) \sum_{i=1}^{k+1} x_i^2 + 2\varepsilon(x_{k+1}) \left(\sum_{i=1}^d x_i \right) \\ &\geq (1 - \alpha) \sum_{i=1}^{k+1} x_i^2 - 2\varepsilon \left| (x_{k+1}) \left(\sum_{i=1}^d x_i \right) \right| \\ &\geq (1 - \alpha) \sum_{i=1}^{k+1} x_i^2 - 2 \left(\frac{1 - \alpha}{k} \right) \left| (x_{k+1}) \left(\sum_{i=1}^d x_i \right) \right| \\ &\geq \left(\frac{1 - \alpha}{k} \right) \left(k \sum_{i=1}^{k+1} x_i^2 - 2 \left| (x_{k+1}) \left(\sum_{i=1}^d x_i \right) \right| \right) \\ &\geq \left(\frac{1 - \alpha}{k} \right) \left(k \sum_{i=1}^{k+1} x_i^2 - \sum_{i=1}^d (x_i^2 + x_{k+1}^2) \right) \\ &\geq 0. \end{aligned} \quad \square$$

Now, we prove the connectivity lemma.

Lemma 5.21. *Fix an integer $k \geq 2$ and $0 < \alpha < 1$. Suppose that U and V are two α -configurations in \mathbb{S}^n with cardinality k each. Then, there exists $n_0 := n_0(k, \alpha)$, and $L := L(k, \alpha)$ such that as long as $n \geq n_0$, there exist α -configurations V_1, V_2, \dots, V_L each with cardinality k such that*

1. The endpoints are U and V i.e., $U = V_1, V = V_L$.
2. Any two consecutive configurations differ in exactly one vector i.e., $|V_i \cap V_{i+1}| = k - 1$ for every $i \in [L-1]$.

Proof. We use induction on k . First, we consider the case when $k = 2$. Let $U = \{\mathbf{u}_1, \mathbf{u}_2\}$, and $V = \{\mathbf{v}_1, \mathbf{v}_2\}$ be two α -configurations. Consider an arbitrary vector \mathbf{w} that is orthogonal to all the vectors in U and V . Such a \mathbf{w} is guaranteed to exist when n is large enough. Now, using Lemma 5.19, we can infer that there exists a configuration $W = \{\mathbf{w}, \mathbf{w}'\}$ such that there is a path of length $O_\alpha(1)$ from U to W , and from V to W . Thus, there exists a path of length $O_\alpha(1)$ 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 \mathbf{w} that is orthogonal to each of the vectors in U and V . Using Lemma 5.19, there are α -configurations $X = \{\mathbf{w}, \mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_{k-1}\}$ and $Y = \{\mathbf{w}, \mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_{k-1}\}$ such that there is a path of length $O_{\alpha,k}(1)$ 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 $O_{\alpha,k}(1)$. We achieve this by restricting ourselves to components orthogonal to \mathbf{w} of the $(k - 1)$ -sized configurations $X' = \{\mathbf{x}'_1, \mathbf{x}'_2, \dots, \mathbf{x}'_{k-1}\}$ and $Y' = \{\mathbf{y}'_1, \mathbf{y}'_2, \dots, \mathbf{y}'_{k-1}\}$, where $\mathbf{x}'_i = \frac{\mathbf{x}_i - \langle \mathbf{x}_i, \mathbf{w} \rangle \mathbf{w}}{\|\mathbf{x}_i - \langle \mathbf{x}_i, \mathbf{w} \rangle \mathbf{w}\|}$ for each $i \in [k - 1]$ (and similarly for \mathbf{y}'_i). Note that X' and Y' are $\frac{\alpha}{1+\alpha}$ -configurations, each of which is a subset of

$$\{\mathbf{v} \in \mathbb{S}^n : \langle \mathbf{v}, \mathbf{w} \rangle = 0\} \cong \mathbb{S}^{n-1}.$$

By the induction hypothesis, there exists a path from X' to Y' using $\frac{\alpha}{1+\alpha}$ -configurations in $\{\mathbf{v} \in \mathbb{S}^n : \langle \mathbf{v}, \mathbf{w} \rangle = 0\}$. Adding the component along \mathbf{w} , we get a path from X to Y of length $O_{\alpha,k}(1)$, finishing the proof. \square

We are now ready to prove Lemma 5.17.

Proof. Suppose for contradiction that there exists a coloring $f : \mathbb{S}^n \rightarrow \{0, 1\}$ that is folded and respects the PCSP Γ_5 . Consider an arbitrary set of vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{2k-1}\}$ that are all orthogonal to \mathbf{v}_0 , and have pairwise inner product α . Such a set is guaranteed to exist as $\alpha \geq 0$. There exists a set of k vectors among these that are all assigned the same color in f . Let these form the configuration U , and the set of negations of these vectors be the configuration V . Using Lemma 5.21, 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 0 and k vectors assigned $+1$ respectively. Since we change at most one vector at a time, there exists a configuration where we have exactly b 1s, a contradiction. \square

Lemmas 5.13 to 5.17 for templates Γ_{1-5} , together with Lemma 5.1, Theorem 5.2 and Lemma 5.6 finish the proof of our main hardness result Theorem 1.4.

Explicit Construction. We give an explicit construction of an integrality gap instance for $\text{fiPCSP}(\Gamma_5)$ where $\Gamma_5 = (P, Q)$, $P = \text{Ham}_k\{1, k\}$, $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ 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 $x_i : i \in [2k - 1], x_S^{(i)} : i \in [L], S \subseteq [2k - 1], |S| = k$. For every subset $S \subseteq [2k - 1]$ with $|S| = k$ and $S = \{i_1, i_2, \dots, i_k\}$, we add the following $L + k + 1$ constraints to the instance:

$$\begin{aligned} &\{x_{i_1}, x_{i_2}, \dots, x_{i_k}\}, \{x_{i_2}, x_{i_3}, \dots, x_{i_k}, x_S^{(1)}\}, \{x_{i_3}, \dots, x_S^{(1)}, x_S^{(2)}\}, \dots, \{x_S^{(L-k+1)}, \dots, x_S^{(L)}\}, \\ &\{x_S^{(L-k+2)}, \dots, x_S^{(L)}, \overline{x_{i_1}}\}, \{x_S^{(L-k+3)}, \dots, \overline{x_{i_1}}, \overline{x_{i_2}}\}, \dots, \{\overline{x_{i_1}}, \overline{x_{i_2}}, \dots, \overline{x_{i_k}}\}. \end{aligned}$$

We choose L to be the constant factor from Lemma 5.21 with $\alpha = \frac{k-3}{k-1}$. The idea is that when all the literals in the constraint $\{x_{i_1}, x_{i_2}, \dots, x_{i_k}\}$ are set to be True, all the literals in $\{\overline{x_{i_1}}, \overline{x_{i_2}}, \dots, \overline{x_{i_k}}\}$ are set to be False. Since there is a series of constraints between them where we alter a single variable, there must exist a constraint where there are exactly b variables that are set to True.

Formally, we show that this instance lacks a weakly satisfying assignment for $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$. Suppose for contradiction that there is an assignment that weakly 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 k of 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 b variables that are set to True, a contradiction.

We now show that the instance has a basic SDP solution with zero error. Set $\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{2k-1} \in \mathbb{S}^n$ to the variables $x_1, x_2, \dots, x_{2k-1}$ such that $\langle \mathbf{v}_i, \mathbf{v}_0 \rangle = 0$, and $\langle \mathbf{v}_i, \mathbf{v}_j \rangle = \alpha$ for every $i \neq j$. We can prove the existence of such a set of vectors by showing that the Gram matrix of the vectors $\{\mathbf{v}_1, \mathbf{v}_2, \dots, \mathbf{v}_{2k-1}\}$ is positive semidefinite.

The Gram matrix G has 1 on the diagonal entries and $\alpha \geq 0$ on the off-diagonal entries:

$$G = \begin{bmatrix} 1 & \alpha & \alpha & \cdots & \alpha \\ \alpha & 1 & \alpha & \cdots & \alpha \\ \alpha & \alpha & 1 & \cdots & \alpha \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ \alpha & \alpha & \alpha & \cdots & 1 \end{bmatrix}.$$

We can write G as:

$$G = \alpha J_{2k-1} + (1 - \alpha)I_{2k-1},$$

where I_{2k-1} is the $(2k - 1) \times (2k - 1)$ identity matrix, and J_{2k-1} is the $(2k - 1) \times (2k - 1)$ matrix with all entries equal to 1.

Since both I_{2k-1} and J_{2k-1} are positive semidefinite, and $0 \leq \alpha \leq 1$, the matrix G is positive semidefinite as well. Finally, we use Lemma 5.21 to set the vectors $\mathbf{v}_S^{(i)}$ for every $S \subseteq [2k - 1], |S| = k, i \in [L]$.

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) [BBKO21].

In this section, we describe the structure of the “basic SDP minion,” which gives a precise algebraic condition for when the basic SDP decides a Promise CSP. We note that a similar minion was concurrently and independently discovered by Ciardo-Zivny [CZ23b].

We prove our minion characterization for “ordinary” promise constraint satisfaction problems over any domain. In particular, we do not consider instances with folding or the setting of constants.

6.1 Minion preliminaries

We now give the definition of a minion. This is most easily stated in the language of category theory, however we shall also give “concrete” interpretations of any category-theoretic statements we make. We let \mathbf{Set} denote the category of sets, and $\mathbf{NeFinSet}$ denote the category of *non-empty* finite sets.

Definition 6.1. *A minion \mathcal{M} is a functor from $\mathbf{NeFinSet}$ to \mathbf{Set} . More concretely, for every non-empty finite set A , we let \mathcal{M}^A denote the image of A with respect to the functor. For every pair of non-empty finite sets A*

and B and function $\pi : A \rightarrow B$, we let the minor map $\mathcal{M}^\pi : \mathcal{M}^A \rightarrow \mathcal{M}^B$ denote the image of π with respect to the functor. Further, for any maps $\pi : A \rightarrow B$ and $\tau : B \rightarrow C$, we have that $\mathcal{M}^\tau \circ \mathcal{M}^\pi = \mathcal{M}^{\tau \circ \pi}$. We also assert for the identity map $\text{id}_A : A \rightarrow A$, we have that $\mathcal{M}^{\text{id}_A} : \mathcal{M}^A \rightarrow \mathcal{M}^A$ is also the identity map.

Remark 6.2. Some works only use minion (or clonoid) to refer to sets of the form $\text{Pol}(\Gamma)$ for some (possibly) promise template Γ , and refer to all other minions as an abstract minions. In this work, all such objects are minions.

In addition, \mathcal{M} is typically only defined as a functor from the subcategory of NeFinSet of sets of the form $[k]$ for some $k \in \mathbb{N}$. Our definition is equivalent as every non-empty finite set A has a bijection with $[|A|]$. However, we consider this richer definition of a minion so that we can discuss how the minion acts on sets without an explicit enumeration.

6.1.1 Examples

We now walk through a few examples.

Polymorphisms. Given a promise template $\Gamma = \{(P_1, Q_1), \dots, (P_l, Q_l)\}$ over the domain (D_1, D_2) , one can view $\text{Pol}(\Gamma)$ as a minion. Formally, for every non-empty finite set A , we let $\text{Pol}(\Gamma)^A$ denote the set of functions $f : D_1^A \rightarrow D_2$ (where D_1^A is the set of functions from A to D_1) such that for any $i \in [l]$ and $\alpha : A \rightarrow P_i$, we have that

$$(f(\pi_1 \circ \alpha), \dots, f(\pi_{k_i} \circ \alpha)) \in Q_i, \quad (9)$$

where $\pi_j : P_i \rightarrow D_1$ is the projection of $P_i \subseteq D_1^{k_i}$ onto the j th coordinate. Observe that if $A = \{1, \dots, L\}$ then $\text{Pol}(\Gamma)^A$ can be identified with the arity L polymorphisms.

Likewise, for any map $\tau : A \rightarrow B$ for non-empty finite sets A and B , we define $\text{Pol}(\Gamma)^\tau$ such that for any $f \in \text{Pol}(\Gamma)^A$ and $\beta : B \rightarrow D_1$, we have that

$$(\tau \otimes f)(\beta) = f(\beta \circ \tau). \quad (10)$$

We It's clear that $\text{id}_A \otimes f = f$. Furthermore, for any pair of maps $\tau : A \rightarrow B$, $\tau' : B \rightarrow C$, and $\gamma : C \rightarrow D_1$, we have that

$$(\tau' \otimes (\tau \otimes f))(\gamma) = (\tau \otimes f)(\gamma \circ \tau') = f((\gamma \circ \tau') \circ \tau) = f(\gamma \circ (\tau' \circ \tau)) = ((\tau' \circ \tau) \otimes f)(\gamma),$$

so associativity also holds.

Trivial minion. The simplest example of a minion $\mathcal{M}_{\text{triv}}$ maps every non-empty finite set to the same set $\{e\}$. All functions $\tau : A \rightarrow B$ map to the same trivial map from e to itself.

Identity minion. The *identity minion* \mathcal{M}_{id} is the identity functor from NeFinSet to itself (when viewed as a subcategory of Set). In other words, for every non-empty finite set A we have that $\mathcal{M}_{\text{id}}^A = A$, and for every map $\tau : A \rightarrow B$, we have that $\mathcal{M}_{\text{id}}^\tau = \tau$.

Convex combinations. For every non-empty finite set A , let Δ_A be the set of probability distributions over A . That is, $\Delta_A = \{f : A \rightarrow [0, 1] : \sum_{a \in A} f(a) = 1\}$. We let $\mathcal{Q}_{\text{conv}}$ denote the minion which maps A to Δ_A [BBKO21]. Given a map $\tau : A \rightarrow B$ and a probability distribution $f \in \Delta_A$, we let $\tau \circ f \in \Delta_B$ be defined by

$$(\tau \circ f)(b) = \sum_{a \in \tau^{-1}(b)} f(a),$$

where the empty sum is equal to 0. The main motivation for studying $\mathcal{Q}_{\text{conv}}$ is that it is closely related to the use of basic LP [BBKO21].

6.1.2 Minion homomorphisms

A key approach of [BBKO21] is that to better understand $\text{PCSP}(\Gamma)$, we can ask how $\text{Pol}(\Gamma)$ relates to other (abstract) minions. The most natural way to do that is via a *minion homomorphism*, which in category-theoretic language is a *natural transformation* between two minions.

Definition 6.3. A minion homomorphism $\psi : \mathcal{M} \rightarrow \mathcal{N}$ between two minions is a natural transformation between the respective functors. More concretely, for every non-empty finite set A , we have a map $\psi_A : \mathcal{M}^A \rightarrow \mathcal{N}^A$ such that for every map $\tau : A \rightarrow B$ and $M \in \mathcal{M}^A$, we have that

$$\tau \circ \psi_A(M) = \psi_B(\tau \circ M).$$

In other words, the following diagram commutes:

$$\begin{array}{ccc} \mathcal{M}^A & \xrightarrow{\psi_A} & \mathcal{N}^A \\ \downarrow \mathcal{M}^\tau & & \downarrow \mathcal{N}^\tau \\ \mathcal{M}^B & \xrightarrow{\psi_B} & \mathcal{N}^B \end{array}$$

To assert the existence of a minion homomorphism from \mathcal{M} and \mathcal{N} , we write $\mathcal{M} \rightarrow \mathcal{N}$. A key result of [BBKO21] is that minion homomorphisms govern the computational complexity of PCSPs.

Theorem 6.4 ([BBKO21]). Let Γ_1, Γ_2 be promise templates. If $\text{Pol}(\Gamma_1) \rightarrow \text{Pol}(\Gamma_2)$ then decision version of $\text{PCSP}(\Gamma_2)$ is polynomial-time reducible to $\text{PCSP}(\Gamma_1)$.

Another result of [BBKO21] is that for any promise template Γ , $\mathcal{Q}_{\text{conv}} \rightarrow \text{Pol}(\Gamma)$ if and only if the basic LP solves $\text{PCSP}(\Gamma)$. Our goal in the remainder of this section is to prove an analogous result for the basic SDP.

6.1.3 The free structure

Let Γ be a promise template. One of the most important tools for understanding minion homomorphisms of the form $\mathcal{M} \rightarrow \text{Pol}(\Gamma)$ is the *free structure* [BBKO21]. Since in this paper we parameterize our promise constraint satisfaction problems using promise templates rather than a pair of relational structures, we restate the definition of a free structure in our language, which we call a *free template*. Given a relation $R \subseteq D^k$, we define the *free relation* $\mathbb{F}_{\mathcal{M}}(R)$ to be a predicate of arity k over \mathcal{M}^D with the following property, we have that $(M_1, \dots, M_k) \in (\mathcal{M}^D)^k$ is an element of R if and only if there exists $N \in \mathcal{M}^R$ such that for all $i \in [k]$, we have that

$$\pi_i \circ N = M_i,$$

where $\pi_i : R \rightarrow D$ is the projection of R onto the i th coordinate.

Given a promise template $\Gamma = \{(P_1, Q_1), \dots, (P_\ell, Q_\ell)\}$, we define the *free relation*

$$\mathbb{F}_{\mathcal{M}}(\Gamma) := \{(\mathbb{F}_{\mathcal{M}}(P_1), Q_1), \dots, (\mathbb{F}_{\mathcal{M}}(P_\ell), Q_\ell)\}.$$

Recall that for a promise template Γ over domain (D_1, D_2) , there exists a map $h : D_1 \rightarrow D_2$ such for every $i \in [\ell]$ and $\mathbf{x} \in P_i$, we have that $h(\mathbf{x}) \in Q_i$. It is not obvious (nor necessarily true) that for $\mathbb{F}_{\mathcal{M}}(\Gamma)$ an analogous map $h^{\mathcal{M}} : \mathcal{M}^{D_1} \rightarrow D_2$ necessarily exists. In fact, constructing $h^{\mathcal{M}}$ asserts the existence of a minion homomorphism. This is the fundamental property of the free structure.

Lemma 6.5 (Lemma 4.4 of [BBKO21]). $\mathcal{M} \rightarrow \text{Pol}(\Gamma)$ if and only if $\mathbb{F}_{\mathcal{M}}(\Gamma)$ is a promise template.

In the language of category theory, the free structure can be thought of as an adjoint functor. See the discussion preceding Remark 4.5 in [BBKO21].

6.2 SDP minion definition

We now define the minion corresponding to the behavior of the basic SDP. We call this minion \mathcal{M}_{SDP} . Recall that for $\mathcal{Q}_{\text{conv}}$, for any non-empty finite set A , we have that $\mathcal{Q}_{\text{conv}}^A = \Delta_A$, where Δ_A is the set of probability distributions over A . We define $\mathcal{M}_{\text{SDP}}^A$ in an analogous manner which captures orthogonal configurations of vectors index by A .

Since the SDP minion is a universal object, we need to be able to represent vectors of arbitrary large dimensions. We achieve this using infinite-dimensional vectors that are eventually zero. Similar techniques have been used in other minion constructions [CZ23a].

Let \mathbb{R}^ω be infinite sequences of real numbers indexed by \mathbb{N} which are eventually 0. We note that \mathbb{R}^ω is an inner product space with the bilinear form

$$\langle \mathbf{x}, \mathbf{y} \rangle := \sum_{i=1}^{\infty} x_i y_i,$$

which is well-defined as only finitely many terms in the summation are nonzero.

We now formally define \mathcal{M}_{SDP} . For each non-empty finite set A , we define $\mathcal{M}_{\text{SDP}}^A$ to be the set of functions $\mathbf{w} : A \rightarrow \mathbb{R}^\omega$ with the following properties.

1. For all distinct $a, a' \in A$, $\langle \mathbf{w}(a), \mathbf{w}(a') \rangle = 0$.
2. $\sum_{a \in A} \|\mathbf{w}(a)\|_2^2 = 1$.

Observe that the first condition implies that the second condition is equivalent to $\|\sum_{a \in A} \mathbf{w}(a)\|_2 = 1$.

The minors of \mathcal{M}_{SDP} are defined analogously to $\mathcal{Q}_{\text{conv}}$. Let A and B be non-empty finite sets and consider $\mathbf{w} \in \mathcal{M}_{\text{SDP}}^A$ and $\tau : A \rightarrow B$. We define $\tau \circ \mathbf{w} : B \rightarrow \mathbb{R}^\omega$ to be for all $b \in B$,

$$(\tau \circ \mathbf{w})(b) := \sum_{a \in \tau^{-1}(b)} \mathbf{w}(a),$$

where the empty sum equals $\mathbf{0} \in \mathbb{R}^\omega$. We now finish the remaining details which prove that \mathcal{M}_{SDP} is a minion.

Lemma 6.6. \mathcal{M}_{SDP} is a minion.

Proof. First, for each $\mathbf{w} \in \mathcal{M}_{\text{SDP}}^A$ and $\tau : A \rightarrow B$ we verify that $\mathbf{w}' := \tau \circ \mathbf{w} \in \mathcal{M}_{\text{SDP}}^B$. First, fix distinct $b, b' \in B$. We have that

$$\langle \mathbf{w}'(b), \mathbf{w}'(b') \rangle = \left\langle \left(\sum_{a \in \tau^{-1}(b)} \mathbf{w}(a) \right), \left(\sum_{a' \in \tau^{-1}(b')} \mathbf{w}(a') \right) \right\rangle = \sum_{\substack{a \in \tau^{-1}(b) \\ a' \in \tau^{-1}(b')}} \langle \mathbf{w}(a), \mathbf{w}(a') \rangle = 0,$$

since b and b' are distinct and τ is a function. Further,

$$\left\langle \left(\sum_{b \in B} \mathbf{w}'(b) \right), \left(\sum_{b \in B} \mathbf{w}'(b) \right) \right\rangle = \left\langle \left(\sum_{a \in A} \mathbf{w}(a) \right), \left(\sum_{a \in A} \mathbf{w}(a) \right) \right\rangle = 1.$$

Thus, $\mathbf{w}' \in \mathcal{M}_{\text{SDP}}^B$. It is trivial to observe that $\mathcal{M}^{\text{id}_A}$ is the identity map.

The only remaining condition to check is that the minors commute. Consider $\pi : A \rightarrow B$ and $\tau : B \rightarrow C$. Consider $\mathbf{u} \in \mathcal{M}_{\text{SDP}}^A$. We seek to verify that $\tau \circ (\pi \circ \mathbf{u}) = (\tau \circ \pi) \circ \mathbf{u}$. For all $c \in C$, we have that

$$(\tau \circ (\pi \circ \mathbf{u}))(c) = \sum_{b \in \tau^{-1}(c)} (\pi \circ \mathbf{u})(b) = \sum_{b \in \tau^{-1}(c)} \sum_{a \in \pi^{-1}(b)} \mathbf{u}(a) = \sum_{a \in (\tau \circ \pi)^{-1}(c)} \mathbf{u}(a) = ((\tau \circ \pi) \circ \mathbf{u})(c),$$

as desired. □

The goal of the rest of this section is to prove the following theorem.

Theorem 6.7. *For any promise template Γ , the basic SDP decides $\text{PCSP}(\Gamma)$ if and only if $\mathcal{M}_{\text{SDP}} \rightarrow \text{Pol}(\Gamma)$.*

6.3 On the basic SDP

As we are considering the decision version of the basic SDP, we assume that $\varepsilon_j = 0$ for each clause index $j \in [m]$. For completeness, we state the simplified SDP below. Note that we omit an objective since we are concerned only with feasibility. Any undefined notation is the same as in Section 2. Given a clause index $j \in [m]$, we let \mathcal{F}_j denote the set of assignments $f : S_j \rightarrow D$ for which $f(C_j) \in P^{(j)}$.

$$\begin{aligned} \text{(a1)} \quad & \lambda_j(f) \geq 0 && \forall j \in [m], f \in \mathcal{F}_j \\ \text{(a2)} \quad & \sum_{f \in \mathcal{F}_j} \lambda_j(f) = 1 && \forall j \in [m] \\ \text{(a3)} \quad & \langle \mathbf{v}_0, \mathbf{v}_0 \rangle = 1 && \\ \text{(a4)} \quad & \langle \mathbf{v}_{i,a}, \mathbf{v}_0 \rangle = \sum_{\substack{f \in \mathcal{F}_j \\ f(x_i)=a}} \lambda_j(f) && \forall j \in [m], x_i \in C_j, a \in D \\ \text{(a5)} \quad & \langle \mathbf{v}_{i,a}, \mathbf{v}_{i',a'} \rangle = \sum_{\substack{f \in \mathcal{F}_j \\ f(x_i)=a, f(x_{i'})=a'}} \lambda_j(f) && \forall j \in [m], x_i, x_{i'} \in C_j, a, a' \in D \end{aligned}$$

We call this formulation of the basic SDP the “traditional” basic SDP, in contrast to the “alternative” basic SDP we soon present.

Remark 6.8. *Although we specify here that the vectors can have an arbitrarily large dimension, it is known that an SDP with n vector or scalar variables is feasible if and only if it is feasible in n dimensions. Thus, nothing is lost from Section 2 by making the vectors have an arbitrarily large (but finite) dimension.*

Remark 6.9. *Some formulations of the basic SDP also impose that for every pair of vectors appearing in the basic SDP has a nonnegative pairwise dot product. Raghavendra’s theorem does not use this in their basic SDP formulation, so we do not include it in our formulation.*

6.3.1 An alternative basic SDP

In order to prove Theorem 6.7, we now present a modified basic SDP which is more convenient to work with. In particular, instead of having an LP variable $\lambda_j(f)$ for each potential clause assignment $f : S_j \rightarrow D$, we instead have a vector $\mathbf{w}_{j,f} \in \mathbb{R}^\omega$. For each $i \in [n]$, let $\mathbf{v}_i : D \rightarrow \mathbb{R}^\omega$ denote the function $\mathbf{v}_i(a) = \mathbf{v}_{i,a}$. Likewise, for every $j \in [m]$, let $\mathbf{w}_j : \mathcal{F}_j \rightarrow \mathbb{R}^\omega$ be the function $\mathbf{w}_j(f) = \mathbf{w}_{j,f}$. With this notation, we can succinctly impose our SDP conditions as follows.

$$\begin{aligned}
\text{(b1)} \quad & \mathbf{v}_i \in \mathcal{M}_{\text{SDP}}^D && \forall i \in [n] \\
\text{(b2)} \quad & \mathbf{w}_j \in \mathcal{M}_{\text{SDP}}^{\mathcal{F}_j} && \forall j \in [m] \\
\text{(b3)} \quad & \mathbf{v}_{i,a} = \sum_{\substack{f \in \mathcal{F}_j \\ f(x_i)=a}} \mathbf{w}_{j,f} && \forall j \in [m], i \in S_j, a \in D,
\end{aligned}$$

where we use the fact that membership in $\mathcal{M}_{\text{SDP}}^S$ can be expressed in $O(|S|^2)$ SDP constraints.

Remark 6.10. *Note that the special “truth” vector \mathbf{v}_0 is omitted from this formulation. As we shall soon see, this does not change the power of the SDP relaxation, but it does make the analysis more straightforward.*

Remark 6.11. *Note that (b3) can also be rewritten as $\mathbf{v}_i = \tau_{i,j} \circledast \mathbf{w}_j$, where $\tau_{i,j} : \mathcal{F}_j \rightarrow D$ is the map $\tau_{i,j}(f) = f(x_i)$. Written in this level of generality, conditions (b1–b3) can then be readily modified to give a relaxation of any minion \mathcal{M} . However, such a relaxation is not algorithmically tractable in general and thus the study of such relaxations in the PCSP literature has been restricted to a restricted class of minions (e.g., “linear” minions in [CŽ23b]).*

6.3.2 Equivalence with traditional basic SDP

We now prove that these two notions formulations of the basic SDP are equivalent. That is, a solution to either SDP can be transferred to the other.

Lemma 6.12. *For any promise template Γ and any instance of PCSP(Γ), the traditional basic SDP for the instance is feasible if and only if the alternative basic SDP for the instance is feasible.*

Before we begin, we note that we often apply linear maps to the vectors into order to transform vectors from one SDP solution to another, extending the ideas in Section 2.3. In particular, we let $\mathbb{R}^{\omega \times \omega}$ denote the set of matrices such that each row and column has finite support. These finite support conditions are such that for any $M \in \mathbb{R}^{\omega \times \omega}$ and $v \in \mathbb{R}^\omega$ we have that $Mv \in \mathbb{R}^\omega$ and that $\mathbb{R}^{\omega \times \omega}$ is closed under the transpose operator.

The identity matrix $I_\omega \in \mathbb{R}^{\omega \times \omega}$ has $(I_\omega)_{i,j} = 1$ if $i = j$ and 0 otherwise. We say that $Q \in \mathbb{R}^{\omega \times \omega}$ is orthogonal if $QQ^T = Q^TQ = I_\omega$. We now extend Proposition 2.17 to our setting.

Proposition 6.13. *Let n be a positive integer and $\mathbf{u}_1, \dots, \mathbf{u}_n, \mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^\omega$ be vectors. We have that $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \langle \mathbf{v}_i, \mathbf{v}_j \rangle$ for all $i, j \in [n]$ if and only if there exists an orthogonal matrix $Q \in \mathbb{R}^{\omega \times \omega}$ such that $\mathbf{v}_i = Q\mathbf{u}_i$ for all $i \in [n]$.*

Proof. Note that for any $\mathbf{u}, \mathbf{u}' \in \mathbb{R}^\omega$ and $Q \in \mathbb{R}^{\omega \times \omega}$ orthogonal, we have that $\langle \mathbf{u}, \mathbf{u}' \rangle = \langle Q\mathbf{u}, Q\mathbf{u}' \rangle$. This proves the “if” direction. We now focus on proving the converse.

Assume that $\langle \mathbf{u}_i, \mathbf{u}_j \rangle = \langle \mathbf{v}_i, \mathbf{v}_j \rangle$ for all $i, j \in [n]$. Let N be a positive integer such that $(\mathbf{u}_i)_j = 0$ and $(\mathbf{v}_i)_j = 0$ for all $i \in [n]$ and $j > N$. We can thus assume that $\mathbf{u}_1, \dots, \mathbf{u}_n, \mathbf{v}_1, \dots, \mathbf{v}_n \in \mathbb{R}^N$. By Proposition 2.17 there exists an orthogonal matrix $Q \in \mathbb{R}^{N \times N}$ such that $Q\mathbf{u}_i = \mathbf{v}_i$ for all $i \in [n]$. We can extend this Q to an orthogonal $Q' \in \mathbb{R}^{\omega \times \omega}$ in the following straightforward manner.

$$Q'_{i,j} = \begin{cases} Q_{i,j} & i, j \leq N \\ 1 & i = j > N \\ 0 & \text{otherwise} \end{cases} \quad \square$$

Alternative to traditional. First, we show that a solution to the alternative basic SDP is a solution to the traditional basic SDP. For each $i \in [n]$, define $\mathbf{v}_{i,D} = \sum_{a \in D} \mathbf{v}_{i,a}$. By (b1), we know that each $\mathbf{v}_{i,D}$ is a unit vector. Further by (b3), we can deduce that $\mathbf{v}_{i,D} = \mathbf{v}_{i',D}$ whenever x_i and $x_{i'}$ appear in a common constraint C_j . Thus, $\mathbf{v}_{i,D} = \mathbf{v}_{i',D}$ whenever x_i and $x_{i'}$ are in the same connected component of the hypergraph with vertices $\{x_1, \dots, x_n\}$ and hyperedges $\{C_1, \dots, C_m\}$. Let $I \subseteq [n]$ and $J \subseteq [m]$ be the indices of the variables and constraints appearing in some connected component. For any orthogonal matrix $Q \in \mathbb{R}^{\omega \times \omega}$, each of the constraints (b1), (b2), and (b3) are preserved when we multiply each vector $\mathbf{v}_{i,a}$ for $i \in I$ and $\mathbf{w}_{j,f}$ for $j \in J$ on the left by Q . Thus, by multiplying by suitable orthogonal matrices (i.e., apply Proposition 6.13 for $n = 1$), we can achieve a solution to (b1-3) such that each $\mathbf{v}_{i,D}$ is the same unit vector for all $i \in [n]$, which we shall call \mathbf{v}_0 . We call this an *aligned* solution to the alternative basic SDP.

We now transform this aligned solution into a solution to the traditional basic SDP. Each $\mathbf{v}_{x,a}$ in the traditional basic SDP is actually the same as it appeared in the aligned solution to the alternative basic SDP. We set \mathbf{v}_0 to the newly-defined \mathbf{v}_0 . Finally, we set $\lambda_j(f) = \|\mathbf{v}_{j,f}\|_2^2$ for all $f \in \mathcal{F}_j$.

To see that the traditional basic SDP is satisfied, note that condition (a1) follows by inspection, and condition (a2) follows from (b2). Conditions (a3) and (a4) follow by substituting $\mathbf{v}_{i,D}$ to \mathbf{v}_0 and applying (b3). Condition (a5) follows by combining (b2) and (b3).

Traditional to alternative. We now show that any solution to the traditional basic SDP is a solution to the alternative basic SDP. Assume we have a traditional basic SDP solution. Recall that $\mathbf{v}_{i,a} \in \mathbb{R}^\omega$ for each $i \in [n]$ and $a \in D$. We shall use the same vectors $\mathbf{v}_{i,a}$ in the solution to the alternative basic SDP. Note that (b1) then follows from (a5) and (a2).

Since there are finitely many vectors, there exists some $N \in \mathbb{N}$ such that each $\mathbf{v}_{i,a}$ (and \mathbf{v}_0) has its support within the first N coordinates. For each $j \in [m]$ and $f \in \mathcal{F}_j$, provisionally assign $\hat{\mathbf{w}}_{j,f}$ to be $\sqrt{\lambda_j(f)} \cdot \mathbf{e}_i$, where $\mathbf{e}_i \in \mathbb{R}^\omega$ is the i th standard unit vector for some $i > N$ chosen uniquely for each constraint-assignment pair. As there are finitely many constraint-assignment pairs, all of these new vectors are in \mathbb{R}^ω .

However, as written the $\hat{\mathbf{w}}_{j,f}$'s are not compatible with the $\mathbf{v}_{i,a}$'s. To rectify this, for each $j \in [m]$, $i \in S_j$, and $a \in D$ we define

$$\hat{\mathbf{v}}_{j,i,a} := \sum_{\substack{j \in \mathcal{F} \\ f(x_i)=a}} \hat{\mathbf{w}}_{j,f}$$

Using (a5), we can readily verify that for any $j \in [m]$, $i, i' \in S_j$, and $a, a' \in D$, we have that $\langle \mathbf{v}_{i,a}, \mathbf{v}_{i',a'} \rangle = \langle \hat{\mathbf{v}}_{j,i,a}, \hat{\mathbf{v}}_{j,i',a'} \rangle$. Thus, by Proposition 6.13, for each $j \in [m]$, there exists an orthogonal matrix $Q_j \in \mathbb{R}^{\omega \times \omega}$ such that $Q_j \hat{\mathbf{v}}_{j,i,a} = \mathbf{v}_{i,a}$ for all $i \in S_j$ and $a \in D$.

For each $j \in [m]$ and $f \in \mathcal{F}_j$, define $\mathbf{w}_{j,f} := Q_j \hat{\mathbf{w}}_{j,f}$. Observe that for each $j \in [m]$, $i \in S_j$, and $a \in D$, we have that

$$\mathbf{v}_{i,a} = Q_j \hat{\mathbf{v}}_{j,i,a} = \sum_{\substack{j \in \mathcal{F} \\ f(x_i)=a}} Q_j \hat{\mathbf{w}}_{j,f} = \sum_{\substack{j \in \mathcal{F} \\ f(x_i)=a}} \mathbf{w}_{j,f},$$

as desired. Thus, conditions (b2) and (b3) hold. Therefore, the two SDP formulations are equivalent, proving Lemma 6.12.

6.4 From minion homomorphism to SDP rounding algorithm

Recall that we are trying to prove that for any promise template Γ , the basic SDP decides $\text{PCSP}(\Gamma)$ if and only if $\mathcal{M}_{\text{SDP}} \rightarrow \text{Pol}(\Gamma)$. We begin by showing the easier direction that the minion homomorphism implies that the basic SDP decides $\text{PCSP}(\Gamma)$. In other words, for any instance of $\text{PCSP}(\Gamma)$, the existence of a vector solution to alternative basic SDP of the instance implies that there exists a weakly satisfying assignment to the instance.

Theorem 6.14. *If $\mathcal{M}_{\text{SDP}} \rightarrow \text{Pol}(\Gamma)$, then the basic SDP decides $\text{PCSP}(\Gamma)$.*

Proof. Assume that (D_1, D_2) is the domain of Γ with homomorphism $h : D_1 \rightarrow D_2$. Consider an instance of $\text{PCSP}(\Gamma)$ on variables $V = \{x_1, \dots, x_n\}$ and clauses $\{C_1, \dots, C_m\}$ such that for all $j \in [m]$, the clause C_j corresponds to $(P^{(j)}, Q^{(j)}) \in \Gamma$ of arity k_j . Let S_j be the set of variables appearing at least once in C_j . We also define \mathcal{F}_j as the set of assignments $f : S_j \rightarrow D_1$ for which $f(C_j) \in P^{(j)}$ and \mathcal{G}_j as the set of assignments $g : S_j \rightarrow D_2$ for which $g(C_j) \in Q^{(j)}$.

By Lemma 6.12, we may without loss of generality assume that there exist vectors $\{\mathbf{v}_{i,a} : i \in [n], a \in D_1\}$ and $\{\mathbf{w}_{j,f} : j \in [m], f \in S_j\}$ such that (b1), (b2), and (b3) hold, where we define $\mathbf{v}_i(a) = \mathbf{v}_{i,a}$ and $\mathbf{w}_j(f) = \mathbf{w}_{j,f}$. If no such vectors exist, we know that our instance is not strongly satisfiable by the contrapositive of Proposition 2.13.

Let $\psi : \mathcal{M}_{\text{SDP}} \rightarrow \text{Pol}(\Gamma)$ be the guaranteed minion homomorphism. Let $\text{id}_{D_1} : D_1 \rightarrow D_1$ be the identity map. We define our weak assignment $\sigma : V \rightarrow D_2$ as follows. For all $i \in [n]$, we set

$$\sigma(x_i) := \psi(\mathbf{v}_i)(\text{id}_{D_1}).$$

Unpacking this, $\psi(\mathbf{v}_i) \in \text{Pol}(A, B)^{D_1}$, so $\psi(\mathbf{v}_i)$ has type $\psi(\mathbf{v}_i) : D_1^{D_1} \rightarrow D_2$. The ‘‘canonical’’ element of $D_1^{D_1}$ is the identity map, so we plug that into the function to get a value in D_2 .

We now seek to show that for every $j \in [m]$, that $\sigma|_{S_j} \in \mathcal{G}_j$. First, we need to eliminate any repetition of variables. To do this, observe that $(\mathcal{F}_j, \mathcal{G}_j)$ can be viewed as a promise relation of arity S_j since it is obtained from $(P^{(j)}, Q^{(j)})$ by setting some tuples in the relation equal. As such, we have that $\text{Pol}(\Gamma) \subseteq \text{Pol}(P^{(j)}, Q^{(j)}) \subseteq \text{Pol}(\mathcal{F}_j, \mathcal{G}_j)$. Therefore, we can think of ψ as a minion homomorphism from \mathcal{M}_{SDP} to $\text{Pol}(\mathcal{F}_j, \mathcal{G}_j)$. Thus, since $\mathbf{w}_j \in \mathcal{M}_{\text{SDP}}^{\mathcal{F}_j}$, we have that $\psi(\mathbf{w}_j) \in \text{Pol}(\mathcal{F}_j, \mathcal{G}_j)^{\mathcal{F}_j}$. Thus, $\psi(\mathbf{w}_j)$ has type $\mathcal{F}_j^{D_1} \rightarrow D_2$.

For $x_i \in S_j$, let $\pi_i : \mathcal{F}_j \rightarrow D_1$ be the evaluation map $\pi_i(f) = f(x_i)$. We check that for all $x_i \in S_j$ and $a \in D_1$, we have that

$$(\pi_i \circ \mathbf{w}_j)(a) = \sum_{\substack{f \in \mathcal{F}_j \\ f(x_i)=a}} \mathbf{w}_{j,f} = \mathbf{v}_i(a), \quad (11)$$

by definition of the alternative basic SDP, so $\pi \circ \mathbf{w}_j = \mathbf{v}_i$. Thus, for all $x_i \in S_j$ we have that

$$\begin{aligned}
\sigma(x_i) &= \psi(\mathbf{v}_i)(\text{id}_{D_1}) && \text{(definition of } \sigma) \\
&= \psi(\pi_i \circ \mathbf{w}_j)(\text{id}_{D_1}) && (11) \\
&= (\pi_i \circ \psi(\mathbf{w}_j))(\text{id}_{D_1}) && (\psi \text{ is minion homomorphism}) \\
&= \psi(\mathbf{w}_j)(\text{id}_{D_1} \circ \pi_i) && (10) \\
&= \psi(\mathbf{w}_j)(\pi_i \circ \text{id}_{\mathcal{F}_j}),
\end{aligned}$$

where for the last line $\text{id}_{D_1} \circ \pi_i = \pi_i = \pi_i \circ \text{id}_{\mathcal{F}_j}$.

Apply (9) with $A = \mathcal{F}_i$ and $(P_i, Q_i) := (\mathcal{F}_j, \mathcal{G}_j)$. Since $\psi(\mathbf{w}_j) \in \text{Pol}(\mathcal{F}_j, \mathcal{G}_j)^{\mathcal{F}_j}$ and $\text{id}_{\mathcal{F}_j} : \mathcal{F}_j \rightarrow \mathcal{F}_j$, we have that

$$(\sigma(x_i) : x_i \in S) = (\psi(\mathbf{w}_j)(\pi_i \circ \text{id}_{\mathcal{F}_j}) : x_i \in S) \in \mathcal{G}_j,$$

as desired. Thus, σ is indeed a satisfying assignment to the original instance. Therefore, the basic SDP decides PCSP(Γ). \square

6.5 From SDP rounding algorithm to minion homomorphism

We now prove the converse. To do this, we adapt the proof techniques of Theorem 4.11 of [CZ23a].

Theorem 6.15. *If the basic SDP decides PCSP(Γ), then $\mathcal{M}_{\text{SDP}} \rightarrow \text{Pol}(\Gamma)$.*

Proof. Assume that (D_1, D_2) is the domain of $\Gamma = \{(P_1, Q_1), \dots, (P_l, Q_l)\}$ with a homomorphism $h : D_1 \rightarrow D_2$ such that $h(P_i) \subseteq Q_i$ for all $i \in [l]$. To simplify the proof, we make the following technical assumption similar to the reduction in Remark 5.3. If a relation $(P, Q) \in \Gamma$ of arity k , then every possible relation which can be created by repeating variables is also in Γ . For example, if P, Q have arity 3, then $(P', Q') \in \Gamma$ where $P' = \{(x, y) \in D_1^2 : (x, y, x) \in P\}$ and $Q' = \{(x, y) \in D_2^2 : (x, y, x) \in Q\}$. Note that $\text{Pol}(P, Q) \subseteq \text{Pol}(P', Q')$, so adding (P', Q') to Γ does not make $\mathcal{M}_{\text{SDP}} \rightarrow \text{Pol}(\Gamma)$ harder to verify. Likewise, since repeating variables is allowed in PCSP instances, whether the basic SDP decides PCSP(Γ) is unaffected by these extra relations.

Assume now that the (alternative) basic SDP decides PCSP(Γ). By Lemma 6.5, it suffices to prove that $\mathbb{F}_{\mathcal{M}_{\text{SDP}}}(\Gamma)$ is a promise structure. In other words, we must construct a map $\hat{h} : \mathcal{M}_{\text{SDP}}^{D_1} \rightarrow D_2$ such that $\hat{h}(\mathbb{F}_{\mathcal{M}_{\text{SDP}}}(P_i)) \subseteq Q_i$ for all $i \in [l]$.

By the De Bruijn-Erdős Theorem¹⁴ [dBE51], it suffices to prove that any finite subsets $F_i \subseteq \mathbb{F}_{\mathcal{M}_{\text{SDP}}}(P_i)$ for $i \in [l]$, we have that there exists $\hat{h} : \mathcal{M}_{\text{SDP}}^{D_1} \rightarrow D_2$ such that $\hat{h}(F_i) \subseteq Q_i$ for all $i \in [l]$.

Note that each $F_i \subseteq (\mathcal{M}_{\text{SDP}}^{D_1})^{k_i}$, let V be the subset of $\mathcal{M}_{\text{SDP}}^{D_1}$ which appears in some tuple of F_i for some $i \in [l]$. It suffices to construct $\hat{h} : V \rightarrow D_2$ rather than all of $\mathcal{M}_{\text{SDP}}^{D_1}$. For each $t \in F := F_1 \cup \dots \cup F_l$, let (P_t, Q_t) be promise pair corresponding to t , and let $S_t \subseteq V$ be the set of distinct elements which appear in t . Since Γ is closed under identification of variables, we can impose on t a clause C_t corresponding to (P'_t, Q'_t) such that each variable of S_t appears exactly once.

We view $\Phi := (V, \{C_t : t \in F\})$ as an instance of PCSP(Γ). If we can show this instance is accepted by the basic SDP, then we know there is an assignment $\hat{h} : V \rightarrow D_2$ which weakly satisfies every clause of $F_1 \cup \dots \cup F_l$ since replacing (P_t, Q_t) with (P'_t, Q'_t) does not change whether the clause $\hat{h}|_{S_t}(t) \in Q_t$ if and only if $\hat{h}|_{S_t}(S_t) \in Q'_t$. Thus, \hat{h} has precisely the property we want.

¹⁴See [CZ23a] or Remark 7.13 of [BBKO21] for additional context.

Thus, we conclude the proof by constructing an alternative basic SDP solution for Φ . For convenience, let $V = \{x_1, \dots, x_n\}$ and for each $t \in F$, let \mathcal{F}_t be the set of possible satisfying assignments $f : S_t \rightarrow D_1$ to C_t .

For each $i \in [n]$, we let $\mathbf{v}_i \in \mathcal{M}_{\text{SDP}}^{D_1}$ be $x_i \in \mathcal{M}_{\text{SDP}}^{D_1}$ itself. This ensures that (b1) is satisfied. For each $t \in F$, let \mathcal{F}_t be the set of satisfying assignments to $f : S_t \rightarrow D_1$ to P'_t . By definition of $\mathbb{F}_{\text{SDP}}(P'_t)$, there exists $\mathbf{w}_t \in \mathcal{M}_{\text{SDP}}^{\mathcal{F}_t}$ such that $\pi_i \circ \mathbf{w}_t = x_i = \mathbf{v}_i$ whenever $x_i \in S_t$, where $\pi_i : \mathcal{F}_t \rightarrow D_1$ is the evaluation map $\pi_i(f) = f(x_i)$. The condition $\mathbf{w}_t \in \mathcal{M}_{\text{SDP}}^{\mathcal{F}_t}$ is precisely (b2) and the equation $\pi_i \circ \mathbf{w}_t = \mathbf{v}_i$ is precisely (b3) (see Remark 6.11). Thus, Φ is indeed accepted by the alternative basic SDP, completing the proof. \square

This completes the proof of Theorem 6.7, giving a minion homomorphism characterization of when the basic SDP works for deciding a PCSP.

7 Conclusion

We studied the robust satisfiability of promise CSPs, and specifically the power of SDPs in this context, revealing a number of new phenomena on both the algorithmic and integrality gap fronts. Our work brings to the fore a number of intriguing questions and directions. We list a few below.

- Can we get a robust fiPCSP dichotomy result for all Boolean symmetric promise templates? We conjecture that every Boolean symmetric folded idempotent PCSP without MAJ or AT polymorphisms does not admit a robust algorithm. In our hardness result Theorem 1.4, we showed the same when there is a single predicate pair. Extending our result to multiple predicate pairs is an interesting challenge.

As a concrete example, consider the promise template $\Gamma = \{(P_1, Q_1), (P_2, Q_2)\}$ where $P_1 = \text{Ham}_k\{\frac{k+1}{2}\}$, $Q_1 = \text{Ham}_k\{0, 1, \dots, k-1\}$, $P_2 = \text{Ham}_t\{1\}$, $Q_2 = \text{Ham}_t\{0, 1, \dots, t-2, t\}$ for odd integers $k, t \geq 3$. We have $\text{AT} \subseteq \text{Pol}(P_1, Q_1)$, $\text{MAJ} \subseteq \text{Pol}(P_2, Q_2)$ while both $\text{AT}, \text{MAJ} \not\subseteq \text{Pol}(\Gamma)$. To obtain the integrality gap for the basic SDP relaxation of fiPCSP(Γ), we need to show that there is no folded sphere coloring that respects Γ . There is a folded sphere coloring $f_1 : \mathbb{S}^n \rightarrow \mathcal{B}$ respecting (P_1, Q_1) , and a folded sphere coloring $f_2 : \mathbb{S}^n \rightarrow \mathcal{B}$ respecting (P_2, Q_2) for every positive integer n . The challenge lies in showing that there is an integer n for which there is no single folded coloring $f : \mathbb{S}^n \rightarrow \mathcal{B}$ that respects both (P_1, Q_1) and (P_2, Q_2) simultaneously. We are able to prove this for small values of k, t under an assumption that Theorem 5.8 extends to the density setting: for any constant $0 < \sigma < 1$, and any tuple S of vectors in \mathbb{S}^n satisfying the conditions in Theorem 5.8, in any subset T of \mathbb{S}^n with spherical measure at least σ , there is a tuple of vectors $S' \subseteq T$ that is congruent to S , as long as $n \geq n_0 := n_0(S, \sigma)$. Both proving the density sphere Ramsey theorem, and showing that Γ does not admit a sphere coloring respecting it are interesting open problems, and the former problem could have applications elsewhere as well.

On the other hand, extending the robust fiPCSP classification to PCSP (i.e., no idempotence or folding) for all Boolean symmetric promise templates would require a better understanding of polymorphisms of arbitrary Boolean symmetric promise templates. We remark that for the decision version of PCSPs for Boolean symmetric promise templates, a dichotomy was first proved for the folded case [BG21], and later, the restriction was removed [FKOS19], where the authors showed that the decision version of a Boolean symmetric PCSP can be solved in polynomial time if and only if it has Threshold, AT or Parity polymorphisms. Our algorithm for the MAJ polymorphisms extends to the setting when

there are threshold polymorphisms, similar to the algorithmic result of [FKOS19]. Combining the polymorphic ideas in [FKOS19] with our sphere coloring results is a potential venue to generalize our hardness results to general Boolean symmetric PCSPs.

Finally, can we get a robust PCSP dichotomy result for general Boolean promise templates? We believe new algorithmic techniques are needed to understand what polymorphic families lead to robust algorithms. For the symmetric folded case, MAJ and AT polymorphisms resulted in robust algorithms while Parity polymorphisms resulted in just the decision version being solved in polynomial time. This suggests that the existence of a suitable notion of noise stable family of polymorphisms could be the key to robust algorithms for PCSPs.

- Can we improve the quantitative aspects of our robust approximation algorithm based on the Majority polymorphism to satisfy a $1 - O(\sqrt{\epsilon})$ fraction of the constraints (which would be optimal [KKMO07] under the UGC)? Even for robust Max-Cut, such an algorithm that makes black-box use of the Majority polymorphism is not known. More generally, we do not know how to translate a polymorphism for the CSP into an approximation or robust satisfaction algorithm for it. For a suitable notion of “approximate” polymorphisms, there is such a connection [BR16]. Of note, in the case of CSP templates of arity k (over an arbitrary domain) with a “near-unanimity operator” (of which MAJ₃ is a special case), a robustness algorithm of $1 - \epsilon$ versus $1 - O(\epsilon^{1/k})$ is known [DKK⁺19].
- For the case of AT polymorphisms, our algorithm incurs an exponential loss, and only satisfies $1 - O\left(\frac{\log \log \frac{1}{\epsilon}}{\log \frac{1}{\epsilon}}\right)$ fraction of the constraints in a $(1 - \epsilon)$ -satisfiable instance. Is this inherent (which is known to be the case for Horn-SAT [GZ12])? Can one show (Unique Games) hardness of how robustly one can approximate concrete promise templates like 1-in-3-SAT vs. NAE-3-SAT?
- For CSPs it is known that if Sherali Adams LP correctly ascertains exact satisfiability, then the CSP has a robust satisfaction algorithm (albeit not LP based). Does this connection remain true for PCSPs (similar to Conjecture 1.5)?
- We have minion characterizations for the solvability of PCSPs by SDPs (this work) and the Sherali Adams hierarchy [CŽ22]. What is the relationship between these minions? Do there exist PCSPs whose exact satisfiability is decided by the basic SDP but which cannot be decided by $O(1)$ levels of the Sherali Adams hierarchy? Note that for CSPs, both these classes are the same and coincide with the class of bounded width CSPs.
- As mentioned earlier, our algorithm for Majority polymorphism can be generalized for arbitrary threshold polymorphisms. Together with the hardness results in [BGS23], this gives a dichotomy w.r.t robust PCSP satisfiability for the class *ordered* Boolean promise templates (whose polymorphisms are monotone functions).

References

- [AB21] Kristina Asimi and Libor Barto. Finitely tractable promise constraint satisfaction problems. In Filippo Bonchi and Simon J. Puglisi, editors, *46th International Symposium on Mathematical Foundations of Computer Science, MFCS 2021, August 23-27, 2021*, volume 202 of *LIPICs*, pages 11:1–11:16, Talinn, Estonia, 2021. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.

- [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, January 5-8, 2020*, pages 1479–1495, Salt Lake City, UT, USA, 2020. SIAM.
- [AD22] Albert Atserias and Víctor Dalmau. Promise constraint satisfaction and width. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1129–1153, Alexandria, Virginia, USA, 2022. SIAM.
- [AGH17] Per Austrin, Venkatesan Guruswami, and Johan Håstad. $(2+\epsilon)$ -SAT is NP-hard. *SIAM J. Comput.*, 46(5):1554–1573, 2017.
- [BBB21] Libor Barto, Diego Battistelli, and Kevin M. Berg. Symmetric promise constraint satisfaction problems: Beyond the boolean case. In *38th International Symposium on Theoretical Aspects of Computer Science, STACS 2021, March 16-19, 2021*, volume 187 of *LIPICs*, pages 10:1–10:16, Saarbrücken, Germany (Virtual Conference), 2021. Schloss Dagstuhl - Leibniz-Zentrum für Informatik.
- [BBKO21] Libor Barto, Jakub Bulín, Andrei A. Krokhin, and Jakub Opršal. Algebraic approach to promise constraint satisfaction. *J. ACM*, 68(4):28:1–28:66, 2021.
- [BG19] Joshua Brakensiek and Venkatesan Guruswami. An algorithmic blend of LPs and ring equations for promise CSPs. In *Proceedings of the Thirtieth Annual ACM-SIAM Symposium on Discrete Algorithms, SODA 2019, January 6-9, 2019*, pages 436–455, San Diego, California, USA, 2019. SIAM.
- [BG20] Joshua Brakensiek and Venkatesan Guruswami. Symmetric polymorphisms and efficient decidability of promise CSPs. In *Proceedings of the 2020 ACM-SIAM Symposium on Discrete Algorithms, SODA 2020, January 5-8, 2020*, pages 297–304, Salt Lake City, UT, USA, 2020. SIAM.
- [BG21] Joshua Brakensiek and Venkatesan Guruswami. Promise constraint satisfaction: Algebraic structure and a symmetric Boolean dichotomy. *SIAM J. Comput.*, 50(6):1663–1700, 2021.
- [BGS23] Joshua Brakensiek, Venkatesan Guruswami, and Sai Sandeep. Conditional dichotomy of Boolean ordered promise CSPs. *TheoretCS*, 2, 2023.
- [BGWŽ20] Joshua Brakensiek, Venkatesan Guruswami, Marcin Wrochna, and Stanislav Živný. The power of the combined basic linear programming and affine relaxation for promise constraint satisfaction problems. *SIAM J. Comput.*, 49(6):1232–1248, 2020.
- [BHPZ21] 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)*, pages 484–503. SIAM, 2021.
- [BJK05] Andrei A. Bulatov, Peter Jeavons, and Andrei A. Krokhin. Classifying the complexity of constraints using finite algebras. *SIAM J. Comput.*, 34(3):720–742, 2005.
- [BK14] Libor Barto and Marcin Kozik. Constraint satisfaction problems solvable by local consistency methods. *J. ACM*, 61(1):3:1–3:19, 2014.

- [BK16] Libor Barto and Marcin Kozik. Robustly solvable constraint satisfaction problems. *SIAM J. Comput.*, 45(4):1646–1669, 2016.
- [BK22] Libor Barto and Marcin Kozik. Combinatorial gap theorem and reductions between promise CSPs. In *Proceedings of the 2022 Annual ACM-SIAM Symposium on Discrete Algorithms (SODA)*, pages 1204–1220, Alexandria, Virginia, USA, 2022. SIAM.
- [BKW17] Libor Barto, Andrei A. Krokhin, and Ross Willard. Polymorphisms, and how to use them. *The Constraint Satisfaction Problem: Complexity and Approximability*, 7:1–44, 2017.
- [BR16] Jonah Brown-Cohen and Prasad Raghavendra. Correlation decay and tractability of CSPs. In Ioannis Chatzigiannakis, Michael Mitzenmacher, Yuval Rabani, and Davide Sangiorgi, editors, *43rd International Colloquium on Automata, Languages, and Programming, ICALP 2016, July 11-15, 2016, Rome, Italy*, volume 55 of *LIPICs*, pages 79:1–79:13. Schloss Dagstuhl - Leibniz-Zentrum für Informatik, 2016.
- [Bul17] Andrei A. Bulatov. A dichotomy theorem for nonuniform CSPs. In *58th IEEE Annual Symposium on Foundations of Computer Science*, pages 319–330, Berkeley, California, USA, 2017. IEEE.
- [BV04] Stephen P. Boyd and Lieven Vandenberghe. *Convex Optimization*. Cambridge University Press, 2004.
- [BWŽ21] Alex Brandts, Marcin Wrochna, and Stanislav Živný. The complexity of promise SAT on non-Boolean domains. *ACM Trans. Comput. Theory*, 13(4):26:1–26:20, 2021.
- [Che09] Hubie Chen. A rendezvous of logic, complexity, and algebra. *ACM Comput. Surv.*, 42(1):2:1–2:32, 2009.
- [CMM09] Moses Charikar, Konstantin Makarychev, and Yury Makarychev. Near-optimal algorithms for maximum constraint satisfaction problems. *ACM Trans. Algorithms*, 5(3):32:1–32:14, 2009.
- [CŽ22] Lorenzo Ciardo and Stanislav Živný. The Sherali-Adams hierarchy for promise CSPs through tensors. *arXiv preprint arXiv:2203.02478*, 2022.
- [CŽ23a] Lorenzo Ciardo and Stanislav Živný. CLAP: A new algorithm for promise CSPs. *SIAM J. Comput.*, 52(1):1–37, 2023.
- [CŽ23b] Lorenzo Ciardo and Stanislav Živný. Hierarchies of minion tests for PCSPs through tensors. In *Proceedings of the 2023 ACM-SIAM Symposium on Discrete Algorithms, SODA 2023, Florence, Italy, January 22-25, 2023*, pages 568–580. SIAM, 2023.
- [dBE51] Nicolaas Govert de Bruijn and Paul Erdős. A colour problem for infinite graphs and a problem in the theory of relations. *Indagationes Mathematicae*, 13:371–373, 1951.
- [DK13] Víctor Dalmau and Andrei A. Krokhin. Robust satisfiability for CSPs: Hardness and algorithmic results. *ACM Trans. Comput. Theory*, 5(4):15:1–15:25, 2013.
- [DKK⁺19] Víctor Dalmau, Marcin Kozik, Andrei Krokhin, Konstantin Makarychev, Yury Makarychev, and Jakub Opršal. Robust Algorithms with Polynomial Loss for Near-Unanimity CSPs. *SIAM Journal on Computing*, 48(6):1763–1795, January 2019.

- [FKOS19] Miron Ficak, Marcin Kozik, Miroslav Olsák, and Szymon Stankiewicz. Dichotomy for symmetric boolean PCSPs. In *46th International Colloquium on Automata, Languages, and Programming, ICALP 2019*, volume 132, pages 57:1–57:12, Patras, Greece, 2019. LIPIcs.
- [Fre04] Robert M Freund. Introduction to semidefinite programming (SDP), 2004.
- [FS01] Uriel Feige and Gideon Schechtman. On the integrality ratio of semidefinite relaxations of MAX CUT. In *Proceedings on 33rd Annual ACM Symposium on Theory of Computing, STOC 2001*, pages 433–442, Crete, Greece, 2001. ACM.
- [GLS93] Martin Grötschel, László Lovász, and Alexander Schrijver. *Geometric Algorithms and Combinatorial Optimization*, volume 2 of *Algorithms and Combinatorics*. Springer Berlin Heidelberg, Berlin, Heidelberg, 1993.
- [GM12] Bernd Gärtner and Jiri Matousek. *Approximation algorithms and semidefinite programming*. Springer Science & Business Media, 2012.
- [GZ12] Venkatesan Guruswami and Yuan Zhou. Tight bounds on the approximability of almost-satisfiable horn SAT and exact hitting set. *Theory Comput.*, 8(1):239–267, 2012.
- [Hås01] Johan Håstad. Some optimal inapproximability results. *J. ACM*, 48(4):798–859, 2001.
- [JCG97] Peter Jeavons, David A. Cohen, and Marc Gyssens. Closure properties of constraints. *J. ACM*, 44(4):527–548, 1997.
- [Jea98] Peter Jeavons. On the algebraic structure of combinatorial problems. *Theor. Comput. Sci.*, 200(1-2):185–204, 1998.
- [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*, pages 23–32, Vancouver, Canada, 2002. IEEE.
- [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*, pages 767–775. ACM, 2002.
- [KKMO07] Subhash Khot, Guy Kindler, Elchanan Mossel, and Ryan O’Donnell. Optimal inapproximability results for MAX-CUT and other 2-variable CSPs? *SIAM J. Comput.*, 37(1):319–357, 2007.
- [KMS18] Subhash Khot, Dor Minzer, and Muli Safra. Pseudorandom sets in Grassmann graph have near-perfect expansion. In *59th IEEE Annual Symposium on Foundations of Computer Science, FOCS 2018*, pages 592–601. IEEE Computer Society, 2018.
- [Lov03] L. Lovász. Semidefinite Programs and Combinatorial Optimization. In Bruce A. Reed and Cláudia L. Sales, editors, *Recent Advances in Algorithms and Combinatorics*, pages 137–194. Springer New York, New York, NY, 2003.
- [MOO10] Elchanan Mossel, Ryan O’Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: invariance and optimality. *Annals of Mathematics*, 171:295–341, 2010.
- [Mos10] Elchanan Mossel. Gaussian bounds for noise correlation of functions. *Geometric and Functional Analysis*, 19:1713–1756, 2010.

- [MR95] Jiří Matoušek and Vojtěch Rödl. On Ramsey sets in spheres. *Journal of Combinatorial Theory, Series A*, 70(1):30–44, 1995.
- [O’D14] Ryan O’Donnell. *Analysis of Boolean Functions*. Cambridge University Press, 2014.
- [O’D16] Ryan O’Donnell. SOS is not obviously automatizable, even approximately. *R. O*, 2016.
- [Rag08] Prasad Raghavendra. Optimal algorithms and inapproximability results for every CSP? In *Proceedings of the 40th Annual ACM Symposium on Theory of Computing, May 17-20, 2008*, pages 245–254, Victoria, British Columbia, Canada, 2008. ACM.
- [Sch78] Thomas J. Schaefer. The complexity of satisfiability problems. In *Proceedings of the 10th Annual ACM Symposium on Theory of Computing, STOC 1978*, pages 216–226, San Diego, California, USA, 1978. ACM.
- [Sch08] Grant Schoenebeck. Linear level Lasserre lower bounds for certain k -CSPs. In *49th Annual IEEE Symposium on Foundations of Computer Science, FOCS 2008, October 25-28, 2008*, pages 593–602, Philadelphia, PA, USA, 2008. IEEE Computer Society.
- [STT07] Grant Schoenebeck, Luca Trevisan, and Madhur Tulsiani. A linear round lower bound for Lovasz-Schrijver SDP relaxations of vertex cover. In *22nd Annual IEEE Conference on Computational Complexity (CCC 2007)*, pages 205–216, San Diego, California, USA, 2007. IEEE Computer Society.
- [Tul09] Madhur Tulsiani. CSP gaps and reductions in the Lasserre hierarchy. In *Proceedings of the 41st Annual ACM Symposium on Theory of Computing, STOC 2009*, pages 303–312, Bethesda, Maryland, USA, 2009. ACM.
- [TŽ17] Johan Thapper and Stanislav Živný. The power of Sherali-Adams relaxations for general-valued CSPs. *SIAM J. Comput.*, 46(4):1241–1279, 2017.
- [TŽ18] Johan Thapper and Stanislav Živný. The limits of SDP relaxations for general-valued CSPs. *ACM Trans. Comput. Theory*, 10(3):12:1–12:22, 2018.
- [Ver18] Roman Vershynin. *High-Dimensional Probability: An Introduction with Applications in Data Science*, volume 47 of *Cambridge Series in Statistical and Probabilistic Mathematics*. Cambridge University Press, 2018.
- [Zhu20] Dmitriy Zhuk. A proof of the CSP dichotomy conjecture. *J. ACM*, 67(5):30:1–30:78, 2020.
- [Zwi98] Uri Zwick. Finding almost-satisfying assignments. In *Proceedings of the Thirtieth Annual ACM Symposium on the Theory of Computing, May 23-26, 1998*, pages 551–560, Dallas, Texas, USA, 1998. ACM.

A Efficient SDP Solution

To begin, state a reformulated basic SDP for which it is easier to reason about approximate solutions. Let $A \bullet X$ denote the entry-wise product of matrices. If \mathcal{F} is a collection of matrices, we let $\text{hull}(\mathcal{F})$ denote the convex hull of these matrices—the set of all convex combinations. If \mathcal{F} is a finite set, then $\text{hull}(\mathcal{F})$ can be described by a finite family of inequalities (and thus with bounded coefficients).

Given a clause C_j with variable set S_j , we let \mathcal{F}_j be the set of square matrices with $|S_j| + 2$ rows and columns with the following properties. For every $\sigma : S_j \rightarrow \mathcal{B}$, we have a matrix $M_\sigma \in \mathcal{F}_j$ corresponding to the outer product of the vector $(1, \sigma, \eta_\sigma)$, where $\eta_\sigma = 0$ if σ is a *satisfying* assignment to C_j and $\eta_\sigma = 1$ otherwise. We also define \mathcal{G}_j to be a set of square matrices with $|S_j| + 1$ rows corresponding to the outer products of the form $(1, \sigma)$ where $\sigma : S_j \rightarrow \mathcal{B}$.

Let X be a symmetric matrix with $n + m + 1$ rows/columns—the last m columns correspond to the “error” in each of the m clauses. Given a set S of indices, we let $X|_S$ denote the square submatrix whose rows and columns come from S .

$$\begin{aligned} \text{minimize: } & \sum_{j=1}^m X_{j+n, j+n} \\ \text{subject to: } & X|_{S_j \cup \{0, j+n\}} \in \text{hull}(\mathcal{F}_j), \forall j \in [m] \\ & X \succeq 0. \end{aligned}$$

By Theorem 5.1.1 of [GM12], for any $\delta = (n/\varepsilon)^{-O(1)}$ we can efficiently find a realization M of X such that M has rational entries, M is PSD, the objective is optimized up to an additive δ , and every linear constraint is satisfied up to an additive δ . For the last condition, this means we “only” have that

$$M|_{S_j \cup \{0, j+n\}} + E_j \in \text{hull}(\mathcal{F}_j),$$

where E_j is an “error” matrix in which all entries have absolute value at most δ . In particular, we only know that $M_{i,i} \in [1 - \delta, 1 + \delta]$ for all $i \in \{0, 1, \dots, n\}$. Let M' be a matrix which is the same as M except $M'_{i,i} = 1 + \delta$ for all $i \in \{0, 1, \dots, n\}$. By definition, $D' := M' - M$ is a diagonal matrix with nonnegative entries, so $M' \succeq 0$. Let $N = \frac{1-\lambda}{1+\delta} \cdot M'|_{\{0,1,\dots,n\}} + \lambda I$ for some $\lambda \approx \Theta(\delta)$. Note that $N \succeq 0$ as well and that $N_{i,i} = 1$ for all $i \in \{0, 1, \dots, n\}$.

We claim that for all $j \in [m]$ that $N|_{\{0\} \cup S_j} \in \text{hull}(\mathcal{G}_j)$. If not, there exist a matrix A and number b such that such that $A \bullet X \leq b$ for all $X \in \text{hull}(\mathcal{G}_j)$ but $A \bullet N > b$. Since $N_{i,i} = 1$ for all $i \in \{0, 1, \dots, n\}$ and $X_{i,i} = 1$ for all $X \in \text{hull}(\mathcal{G}_j)$, we may assume without loss of generality that $A_{i,i} = 0$ for all $i \in \{0, 1, \dots, n\}$. Likewise, since all matrices in $\text{hull}(\mathcal{G}_j)$ are symmetric, we may assume that A is symmetric. Thus, since $I \in \text{hull}(\mathcal{G}_j)$, we may assume that $b \geq 0$. In fact, we now show for nonzero A we need $b > 0$.

Claim A.1. *Let A be a symmetric matrix with an all-zero diagonal such that $A \bullet X \leq 0$ for all $X \in \text{hull}(\mathcal{G}_j)$. Then, A is the zero matrix.*

Proof. Consider indices $i < j$ and let $\mathcal{G}'_j \subseteq \mathcal{G}_j$ be the matrices $Y \in \mathcal{G}_j$ for which $Y_{i,j} = 1$. Let Z be the (uniform) average of all matrices in \mathcal{G}'_j . It is straightforward to verify that $Z_{i,j} = Z_{j,i} = 1$, but all other off-diagonal entries are zero. Since $Z, -Z \in \text{hull}(\mathcal{G}_j)$, we have that

$$\begin{aligned} A_{i,j} &= A \bullet Z \leq 0, \\ -A_{i,j} &= A \bullet (-Z) \leq 0. \end{aligned}$$

Therefore, $A_{i,j} = 0$ for all $i < j$. Thus, A is the zero matrix. \square

Since $\text{hull}(\mathcal{G}_j)$ is a projection of $\text{hull}(\mathcal{F}_j)$, we have that this inequality is approximately satisfied by M :

$A \bullet M|_{\{0\} \cup S_j} \leq b + O(\delta)$. Thus,

$$\begin{aligned}
b &< A \bullet N \\
&= A \bullet \left(\frac{1-\lambda}{1+\delta} \cdot (M+D')|_{\{0\} \cup S_j} + \lambda I \right) \\
&= \frac{1-\lambda}{1+\delta} \cdot (A \bullet M|_{\{0\} \cup S_j} + \frac{1-\lambda}{1+\delta} \cdot (A \bullet D'|_{\{0\} \cup S_j} + \lambda(A \bullet I)) \\
&\leq \frac{1-\lambda}{1+\delta} b + O(\delta),
\end{aligned}$$

where we use the fact that every entry of D' is bounded in absolute value by δ and $A \bullet I = 0$. Thus, for $\lambda = \Omega(\delta)$, we get a contradiction. This proves that $N|_{\{0\} \cup S_j} \in \text{hull}(\mathcal{G}_j)$ for all $j \in [m]$.

Using Gaussian elimination (see O'Donnell [O'D16] and references therein, including [Lov03]), we can efficiently compute rational square matrices L and D , where L is lower triangular and D is a diagonal, such that $N = LDL^T$. Then, we can compute vectors $\mathbf{v}_0, \dots, \mathbf{v}_n$ whose coordinates are square roots of rational numbers such that $N_{i,j} = \langle \mathbf{v}_i, \mathbf{v}_j \rangle$. Since $N|_{\{0\} \cup S_j} \in \text{hull}(\mathcal{G}_j)$ for all $j \in [m]$, by solving a suitable linear program, we can efficiently find a probability distribution λ_j over assignments $S_j \rightarrow D$ which is consistent with the first and second moment conditions of the fiPCSP basic SDP, where the objective value is to minimize the total weight ε_j on non-assignments to C_j .

To finish, it suffices to prove that $\varepsilon_j \leq M_{j+n,j+n} + O(\delta)$ for all $j \in [m]$. Recall that $M|_{S_j \cup \{0,j+n\}} + E_j \in \text{hull}(\mathcal{F}_j)$. Thus, $N' := (M|_{S_j \cup \{0,j+n\}} + E_j)|_{\{0\} \cup S_j} \in \text{hull}(\mathcal{G}_j)$. Furthermore, if we solve the same ‘‘error’’ LP for N' , we will get an error ε'_j which must be at most $M_{j+n,j+n} + \delta$. One can see that every entry in N and N' differ by at most $O(\delta)$. Let $\varepsilon : \text{hull}(\mathcal{G}_j) \rightarrow [0, 1]$ be the value of the LP on all possible input matrices. It is clear that ε is a bounded convex function, so it must be Lipschitz. Thus, $\varepsilon_j - \varepsilon'_j \leq O(\delta)$, so $\varepsilon_j \leq M_{j+n,j+n} + O(\delta)$, as desired.

B Missing Proofs

B.1 Missing proofs from Section 4

Proof of Lemma 4.11. We first consider the case when $k = 1$. Without loss of generality, let $P = Q = \{+1\}$, and we use $\mathbf{v}_1 = \alpha \mathbf{v}_0 + \mathbf{v}'_1$ to denote the SDP vector corresponding to the variable used in the constraint. As the basic SDP has error at most $\sqrt{\varepsilon}$, we get that

$$\alpha_1 \geq 1 - \sqrt{\varepsilon}$$

As $\alpha_1^2 + \|\mathbf{v}'_1\|_2^2 = 1$, $\|\mathbf{v}'_1\|_2 \leq O(\varepsilon^{0.25})$. Thus, using Proposition 2.21, we get that $\langle \zeta, \mathbf{v}'_1 \rangle \leq O(\varepsilon^{0.25}r)$ with probability at least $1 - e^{-\frac{1}{2}} \geq 1 - \sqrt{\varepsilon}$. On the other hand, using Proposition 2.20, we get that $|\langle \zeta, \mathbf{v}_0 \rangle| \geq \frac{1}{r}$ with probability at least $1 - \frac{1}{r}$. This implies that

$$\delta \alpha_1 |\langle \zeta, \mathbf{v}_0 \rangle| \geq \frac{\delta}{2r}$$

Thus, with probability at least $1 - O(\frac{1}{r})$, we have

$$\langle \zeta, \mathbf{v}'_1 \rangle \leq O(\varepsilon^{0.25}r) < \frac{\delta}{2r} \leq \delta \alpha_1 |\langle \zeta, \mathbf{v}_0 \rangle|$$

Hence, with probability at least $1 - O(\frac{1}{r})$, we round the variable to $+1$.

We now consider the general case when $k \geq 2$. Note that the above proof for $k = 1$ holds when $P = \text{Ham}_k\{0\}$ or when $P = \text{Ham}_k\{k\}$. We are left with the setting when $P = \text{Ham}_k\{0, k\}$. In order to show that our algorithm is a robust algorithm for this fiPCSP, it suffices to show that all the elements in the predicates are rounded to the same value with high probability. Consider $i, j \in [k]$. 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 a probability of at least $1 - c$ on P .

1. (First moment.) We have

$$|\mu_i - \mu_j| \leq 2c.$$

2. (Second moment.) We have

$$\langle \mathbf{v}_i, \mathbf{v}_j \rangle \geq 1 - 2c$$

Using this, we get

$$\begin{aligned} \|\mathbf{v}'_i - \mathbf{v}'_j\|_2^2 &= \mathbf{v}_i - \mathbf{v}_j + (\alpha_j - \alpha_i)\mathbf{v}_0 \\ &\leq \|\mathbf{v}_i - \mathbf{v}_j\|_2^2 + (\alpha_i - \alpha_j)^2 \\ &\leq O(c). \end{aligned}$$

$$\text{Thus, } \|\mathbf{v}'_i - \mathbf{v}'_j\| \leq O(\sqrt{c}).$$

As earlier, we assume that c is at most $\sqrt{\varepsilon}$.

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

$$\begin{aligned} \langle \zeta, \mathbf{v}'_i \rangle &\geq \delta \alpha_i |\langle \zeta, \mathbf{v}_0 \rangle| \\ \langle \zeta, \mathbf{v}'_j \rangle &< \delta \alpha_j |\langle \zeta, \mathbf{v}_0 \rangle| \end{aligned}$$

Using Proposition 2.21, we can infer that

$$|\langle \zeta, \mathbf{v}'_i \rangle - \langle \zeta, \mathbf{v}'_j \rangle| \leq O(r\sqrt{c})$$

with probability at least $1 - \frac{1}{r}$.

We consider two cases: first, when $|\alpha_i| \leq \frac{1}{2}$. As $\alpha_i^2 + \|\mathbf{v}'_i\|_2^2 = 1$, and $|\alpha_i - \alpha_j| \leq 2c$, we get that $\|\mathbf{v}'_j\| = \Omega(1)$. In this case, we have

$$\begin{aligned} \langle \zeta, \mathbf{v}'_j \rangle &\geq \langle \zeta, \mathbf{v}'_i \rangle - O(r\sqrt{c}) \\ &\geq \delta \alpha_i |\langle \zeta, \mathbf{v}_0 \rangle| - O(r\sqrt{c}) \end{aligned}$$

Thus, we have

$$\langle \zeta, \mathbf{v}'_j \rangle \in [\delta \alpha_i |\langle \zeta, \mathbf{v}_0 \rangle| - O(r\sqrt{c}), \delta \alpha_j |\langle \zeta, \mathbf{v}_0 \rangle|]$$

Here, $\langle \zeta, \mathbf{v}'_j \rangle \in [p, q]$ where $q - p \leq O(\delta r) + O(r\sqrt{c}) \leq O(\delta r)$. However, as $\|\mathbf{v}'_j\| \geq \Omega(1)$, this happens with probability at most $O(\delta r)$.

Now, suppose that $|\alpha_i| \geq \frac{1}{2}$. We have

$$\delta\alpha_i |\langle \zeta, \mathbf{v}_0 \rangle| \geq \delta\alpha_j |\langle \zeta, \mathbf{v}_0 \rangle| - 2\delta c |\langle \zeta, \mathbf{v}_0 \rangle|$$

However, as $\langle \zeta, \mathbf{v}_0 \rangle \sim \mathcal{N}(0, 1)$, we have $|\langle \zeta, \mathbf{v}_0 \rangle| \leq r$ with probability at least $1 - \sqrt{\varepsilon}$. Thus, with probability at least $1 - \sqrt{\varepsilon}$, we have

$$\langle \zeta, \mathbf{v}'_i \rangle \geq \delta\alpha_i |\langle \zeta, \mathbf{v}_0 \rangle| \geq \langle \zeta, \mathbf{v}'_j \rangle - 2\delta cr$$

We have $\delta\alpha_i |\langle \zeta, \mathbf{v}_0 \rangle| \in [p, q]$ where $q - p \leq O(\delta cr) + O(r\sqrt{c})$. However, this happens with probability at most $O(\frac{r\sqrt{c}}{\delta}) \leq O(\frac{1}{r})$. \square

Proof of Lemma 4.12. Suppose that $\mathbf{a} = \text{sgn}(\mathbf{x} - \mathbf{y})$ for $\mathbf{x}, \mathbf{y} \in \text{Aff}(P)$, and $x_i \neq y_i$ for all $i \in [k]$. By modifying the affine combinations slightly, we can assume that \mathbf{x} and \mathbf{y} are rational affine combinations of P while still preserving the fact that $\mathbf{a} = \text{sgn}(\mathbf{x} - \mathbf{y})$. In other words, there exist $p_1, p_2, \dots, p_K, q_1, q_2, \dots, q_K \in \mathbb{Q}$ such that $\sum_{i \in [K]} p_i = \sum_{i \in [K]} q_i = 1$, and $\mathbf{x} = \sum_{i \in [K]} p_i \mathbf{a}_i$, $\mathbf{y} = \sum_{i \in [K]} q_i \mathbf{a}_i$, where $\mathcal{B}^k = \{\mathbf{a}_1, \mathbf{a}_2, \dots, \mathbf{a}_K\}$. Let N be a positive integer such that we can write $p_i = \frac{p'_i}{N}$, $q_i = \frac{q'_i}{N}$ where p'_i, q'_i are integers for every $i \in [K]$.

Let S be a list of elements of \mathcal{B}^k where for each $i \in [K]$, there are $|p'_i| + |q'_i|$ copies of \mathbf{a}_i . Let $\mathbf{s} \in \mathcal{B}^{|S|}$ be a ± 1 vector such that for each $i \in [K]$ we have that $|p'_i|$ of the $j \in [|S|]$ with $S_j = \mathbf{a}_i$ have $s_j = \text{sgn}(p'_i)$ and the other $|q'_i|$ copies have $s_j = \text{sgn}(q'_i)$. Since $\sum_{i \in [K]} p'_i = \sum_{i \in [K]} q'_i$, we have that s_j has an equal number of each of $+1$ and -1 .

Let \mathbf{z} denote the signed sum of all vectors (including repetitions) in S according to the signs of \mathbf{s} . Note that $\text{sgn}(\mathbf{z}) = \text{sgn}(\mathbf{x} - \mathbf{y})$. As each element of \mathbf{z} is an integer, we get that the absolute value of each coordinate in \mathbf{z} is at least 1. Furthermore, we can take multiple copies of S to ensure that the absolute value of each coordinate in \mathbf{z} is at least 2. Now, we add an arbitrary element of P with sign $+1$ to S . Note that we still have that the signed sum of S , i.e., the updated \mathbf{z} satisfies $\text{sgn}(\mathbf{z}) = \text{sgn}(\mathbf{x} - \mathbf{y})$. Furthermore, $\mathbf{z} = \mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L$ where each $\mathbf{x}_i \in P$. Thus, $\mathbf{a} = \text{sgn}(\mathbf{x} - \mathbf{y}) = \text{sgn}(\mathbf{z}) = \text{sgn}(\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L) \in O_{AT}(P)$. Thus,

$$\{\text{sgn}(\mathbf{x} - \mathbf{y}) : \mathbf{x}, \mathbf{y} \in \text{Aff}(P), \forall i, x_i \neq y_i\} \subseteq O_{AT}(P)$$

To prove the other direction, suppose that $\mathbf{a} \in O_{AT}(P)$. That is, $\mathbf{a} = \text{sgn}(\mathbf{x}_1 - \mathbf{x}_2 + \dots + \mathbf{x}_L)$. Let S be a multiset of $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_L$ with the corresponding sign as in the summation. As P is non-trivial in every coordinate i.e., for every $i \in [k]$, there exist assignments \mathbf{x} in P where $x_i = +1$, and similarly, $\mathbf{x}' \in P$ 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 $\text{sgn}(\mathbf{a})$. We modify S while still preserving this property to ensure that the signed sum of vectors in S has an absolute value of at least 2 in every coordinate.

As there are an odd number of vectors in S , the signed sum of the vectors has an absolute value of at least 1 in every coordinate. Fix a vector $\mathbf{x}_i \in S$. 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 $2k$ times to ensure that in the final multiset S , the signed sum has an absolute value of at least 2 in every coordinate. Finally, we add an arbitrary vector with sign -1 to S , to ensure that the number of vectors with $+1$ sign and the number of vectors with -1 sign are equal in S . Overall, we get that there are $\mathbf{x}_1, \mathbf{x}_2, \dots, \mathbf{x}_N \in P$ and $\mathbf{y}_1, \mathbf{y}_2, \dots, \mathbf{y}_N \in P$ such that

$$\mathbf{a} = \text{sgn}(\mathbf{x}_1 + \dots + \mathbf{x}_N - \mathbf{y}_1 - \dots - \mathbf{y}_N) = \text{sgn}\left(\frac{1}{N}\mathbf{x}_1 + \dots + \frac{1}{N}\mathbf{x}_N - \frac{1}{N}\mathbf{y}_1 - \dots - \frac{1}{N}\mathbf{y}_N\right)$$

Thus, we get that $\mathbf{a} \in \{\text{sgn}(\mathbf{x} - \mathbf{y}) : \mathbf{x}, \mathbf{y} \in \text{Aff}(P), \forall i, x_i \neq y_i\}$, completing the proof that

$$O_{AT}(P) \subseteq \{\text{sgn}(\mathbf{x} - \mathbf{y}) : \mathbf{x}, \mathbf{y} \in \text{Aff}(P), \forall i, x_i \neq y_i\} \quad \square$$

B.2 Missing proofs from Section 5

Proof of Lemma 5.1. We extensively use the properties of AT, MAJ polymorphisms of Boolean symmetric folded idempotent PCSPs proved in [BG21].¹⁵ We recall that $O_{AT}(P)$ (respectively $O_{MAJ}(P)$) denotes the set $\bigcup_{\mathbf{x}_1, \dots, \mathbf{x}_L \in P, L \in \mathbb{N}, \text{ odd}} AT_L(\mathbf{x}_1, \dots, \mathbf{x}_L)$ (respectively MAJ_L) for a predicate P .

Let k denote the arity of P, Q , i.e., $P \subseteq Q \subseteq \mathcal{B}^k$. Note that $P \not\subseteq \{(-1, \dots, -1), (+1, \dots, +1)\}$ as in that case $O_{MAJ}(P) = P \subseteq Q$, contradicting the fact that $MAJ \not\subseteq \text{Pol}(P, Q)$. Thus, there exists $l \in \{1, 2, \dots, k-1\}$ such that $\text{Ham}_k\{l\} \subseteq P$.

Case 1. We first consider the case when $P = \text{Ham}_k\{l\}$ for some $l \in \{1, 2, \dots, k-1\}$. As P is symmetric, $O_{MAJ}(P)$ is symmetric as well [BG21]. Furthermore, as $MAJ \not\subseteq \text{Pol}(P, Q)$, there exists $b \in \{0, 1, \dots, k\}$ such that $\text{Ham}_k\{b\} \cap Q = \emptyset$ and $\text{Ham}_k\{b\} \subseteq O_{MAJ}(P)$.

Suppose that $b \notin \{0, k\}$. Let $Q' = \mathcal{B}^k \setminus \text{Ham}_k\{b\}$. By definition, $MAJ \not\subseteq \text{Pol}(P, Q')$. Using the fact that $O_{AT}(\text{Ham}_k\{l\}) = \text{Ham}_k\{1, 2, \dots, k-1\}$, we get that $AT \not\subseteq \text{Pol}(P, Q')$. Thus, we get a promise template (P, Q) that is fipp-definable from original promise template where $P = \text{Ham}_k\{l\}$, $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \in \{1, 2, \dots, k-1\} \setminus \{l\}$. Note that $MAJ, AT \notin \text{Pol}(P, Q)$. We now modify this promise template further, updating P, Q, l, k, b while preserving the following two properties:

1. At every step, $P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \in \{1, 2, \dots, k-1\} \setminus \{l\}$.
2. $MAJ, AT \notin \text{Pol}(P, Q)$.

As $O_{MAJ}(P) = \text{Ham}_k\{0, 1, \dots, k\} \cap \{2l - k + 1, \dots, 2l - 1\}$, and $\text{Ham}_k\{b\} \in O_{MAJ}(P)$, we get that $b \in \{2l - k + 1, \dots, 2l - 1\} \cap \{0, \dots, k\}$. Furthermore, as $b > 0$, we get that $l > 1$. Similarly, we get that $l < k - 1$. This also implies that $k \geq 4$ as $l \in \{1, \dots, k-1\}$.

We use the following two tools to repeatedly obtain a new promise template that is fipp-definable from the previous one.

1. Given a promise template $P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \in \{1, 2, \dots, k-1\} \setminus \{l\}$, then the promise template $P' = \text{Ham}_{k-1}\{l\}, Q' = \text{Ham}_{k-1}\{0, 1, \dots, k-1\} \setminus \{b\}$ is fipp-definable from (P, Q) (Claim 4.2 of [BG21]). As long as $b < k - 1$ and $b \neq 2l - k + 1$, this update preserves the above two properties.
2. Given a promise template $P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \in \{1, 2, \dots, k-1\} \setminus \{l\}$, then the promise template $P' = \text{Ham}_{k-1}\{l-1\}, Q' = \text{Ham}_{k-1}\{0, 1, \dots, k-1\} \setminus \{b-1\}$ is fipp-definable from (P, Q) (Claim 4.4 of [BG21]). As long as $b > 1$ and $b \neq 2l - 1$, this update preserves the above two properties.

Now, we update the promise template using the above two steps. As k is decreasing at every step, this procedure terminates at some point. Then, either of the two conditions holds:

1. $b = 1, b = 2l - k + 1$. In this case, we get that $l = \frac{k}{2}$ and $b = 1$. Thus, (P', Q') is fipp-definable from Γ where $P' = \text{Ham}_k\{\frac{k}{2}\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{1\}$ where k is even and is at least 4.

¹⁵The notion of ppp-reduction used in [BG21] implicitly allows for equality constraints. However, the ppp-reductions we use in this proof (particularly those in Claims 4.2 and 4.4 of [BG21]) do not use equalities.

2. $b = k - 1, b = 2l - 1$. In this case, we get that $l = \frac{k}{2}$ and $b = k - 1$. Thus, (P', Q') is fipp-definable from Γ where $P' = \text{Ham}_k\{\frac{k}{2}\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{k - 1\}$ where k is even and is at least 4.

Suppose that there is no $b \notin \{0, k\}$ such that $\text{Ham}_k\{b\} \subseteq O_{\text{MAJ}}(P) \setminus Q$. As $O_{\text{MAJ}}(P) \not\subseteq Q$, by negating the variables if needed, we can assume that $\text{Ham}_k\{0\} \in O_{\text{MAJ}}(P) \setminus Q$. Furthermore, there exists $b \in \{1, 2, \dots, k - 1\}$ such that $\text{Ham}_k \not\subseteq Q$ as $O_{\text{AT}}(P) \not\subseteq Q$. Thus, we obtain (P, Q) that is fipp-definable from the original promise template such that $P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{1, \dots, k\} \setminus \{b\}$ where $l, b \in \{1, 2, \dots, k - 1\}, b > 2l - 1$. By using the first type of update used above (Claim 4.2 of [BG21]), we obtain a new promise template (P, Q) that is fipp-definable from the original promise template such that $P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{0, k - 1\}$, where $l \in \{1, 2, \dots, k - 1\}, l \leq \frac{k-1}{2}$.

Case 2. There exist $l \neq l'$ such that $\text{Ham}_k\{l, l'\} \subseteq P$. Recall that $P \not\subseteq \text{Ham}_k\{0, k\}$. This implies that $O_{\text{AT}}(P) = \text{Ham}_k\{0, 1, \dots, k\}$. Hence, we can get a promise template (P, Q') that is fipp-definable from the original promise template such that $Q' = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ and $\text{Ham}_k\{b\} \not\subseteq O_{\text{MAJ}}(P)$.

For ease of notation, let $P = \text{Ham}_k S$, where $S \subseteq \{0, 1, \dots, k\}$. First, consider the case when $\min S = 0, \max S = k$. As mentioned earlier, we know that there exists $l \in \{1, 2, \dots, k - 1\}$ such that $\text{Ham}_k\{l\} \subseteq P$. Thus, we can obtain a new PCSP (P, Q) that is fipp-definable from the original promise template where $P = \text{Ham}_k\{0, l, k\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ and $b \notin \{0, l, k\}$. We consider three cases separately:

1. Suppose that $l \leq \frac{k-1}{2}$. In this case, we have a promise template (P, Q) that is fipp-definable from the original promise template where $P = \text{Ham}_k\{l, k\}$ and $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ and $b \notin \{l, k\}$. Note that this promise template does not contain AT or MAJ as polymorphisms.
2. Suppose that $l = \frac{k}{2}$. In this case, we have a promise template (P, Q) that is fipp-definable from the original promise template where $P = \text{Ham}_k\{l\}$ and $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \notin \{0, l, k\}$. This also doesn't have AT and MAJ as polymorphisms. We have already shown that we can relax this further to the earlier mentioned three promise templates.
3. Suppose that $l \geq \frac{k+1}{2}$. In this case, we have a promise template (P, Q) that is fipp-definable from the original PCSP where $P = \text{Ham}_k\{0, l\}$ and $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \notin \{0, l\}$. Note that this promise template does not contain AT or MAJ as polymorphisms.

Thus, there is a promise template (P, Q) that is fipp-definable from the original PCSP where $P = \text{Ham}_k\{l, l'\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ such that $l < l', \{l, l'\} \neq \{0, k\}, b \in O_{\text{MAJ}}(P)$. We end up with the same PCSP when $\{\min S, \max S\} \neq \{0, k\}$.

If $\{l, l'\} = \{1, k\}$ or $\{0, k - 1\}$, we get that the promise template $P = \text{Ham}_k\{1, k\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ is fipp-definable from the original promise template, and we are done. If not, we update the promise template while maintaining the two below properties:

1. $P = \text{Ham}_k\{l, l'\}$ with $l < l'$ and $\{l, l'\} \neq \{0, k\}$ and $Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$. We also assume that $\{l, l'\} \neq \{1, k\}$ and $\{l, l'\} \neq \{0, k - 1\}$.
2. $\text{AT}, \text{MAJ} \not\subseteq \text{Pol}(P, Q)$.

As with the earlier case, we obtain a series of new promise templates that are fipp-definable from the previous promise templates using the below tools.

1. We get $P' = \text{Ham}_{k-1}\{l, l'\}$ and $Q = \text{Ham}_{k-1}\{0, 1, \dots, k\} \setminus \{b\}$ using Claim 4.2 of [BG21]. For this to preserve the above properties, we need that $l' \neq k, b \neq k$ and $b \neq 2l - k + 1$.

2. We get $P' = \text{Ham}_{k-1}\{l-1, l'-1\}$ and $Q = \text{Ham}_{k-1}\{0, 1, \dots, k\} \setminus \{b-1\}$ using Claim 4.4 of [BG21]. For this to preserve the above properties, we need that $l \neq 0, b \neq 0$ and $b \neq 2l' - 1$.

As the arity of the predicates is decreasing at each step, this process terminates in finite steps. When we are unable to update the promise template using the above procedures, one of the following must be true.

1. $l' = k, b = 0$. In this case, we have a promise template (P, Q) where $P = \text{Ham}_k\{l, k\}, Q = \text{Ham}_k\{1, 2, \dots, k\}$, where $l \neq 0, l \leq \frac{k-1}{2}$.
2. $b = k, l = 0$. By negating the variables, we can observe that the above promise template is fipp-definable from this promise template.
3. $b = k, b = 2l' - 1$. We have $l' = \frac{k+1}{2}$. In this case, we have promise template (P, Q) where $P = \text{Ham}_k\{l, \frac{k+1}{2}\}, Q = \text{Ham}_k\{0, 1, 2, \dots, k-1\}$, where $l \leq \frac{k-1}{2}$.
4. We have $b = 2l - k + 1, b = 0$. By negating the variables, we can observe that the above promise template is fipp-definable from this promise template.

Thus, we have obtained a new promise template that is fipp-definable from the original promise template and is equal to either of the following promise templates.

1. k is even, and $P = \text{Ham}_k\{\frac{k}{2}\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ where $b \in \{1, k-1\}$.
2. k is odd, $P = \text{Ham}_k\{l, \frac{k+1}{2}\}, Q = \text{Ham}_k\{0, 1, 2, \dots, k-1\}$, where $l \leq \frac{k-1}{2}$.
3. $P = \text{Ham}_k\{l, k\}, Q = \text{Ham}_k\{1, 2, \dots, k\}$, where $l \neq 0, l \leq \frac{k-1}{2}$.
4. $P = \text{Ham}_k\{l\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{0, k-1\}$ where $l \in \{1, 2, \dots, k-1\}, l \leq \frac{k-1}{2}$.
5. $P = \text{Ham}_k\{1, k\}, Q = \text{Ham}_k\{0, 1, \dots, k\} \setminus \{b\}$ for arbitrary b .

Finally, we note that if a Boolean promise template $\Gamma' = \{(P_1, Q_1), \dots, (P_l, Q_l)\}$ is fipp-definable from another Boolean promise template Γ , then Γ' remains fipp-definable from Γ even when we disallow the constraints of Γ' in the fipp-definition to use constants. \square