Approximating Constraint Satisfaction Problems on High-Dimensional Expanders

Vedat Levi Alev
Computer Science Department
University of Waterloo
Waterloo, Canada
vlalev@uwaterloo.ca

Fernando Granha Jeronimo
Computer Science Department
University of Chicago
Chicago, USA
granha@uchicago.edu

Madhur Tulsiani
Toyota Technological Institute
Chicago, USA
madhurt@ttic.edu

***Abstract*—We consider the problem of approximately solving constraint satisfaction problems with arity $k > 2$ (k-CSPs) on instances satisfying certain expansion properties, when viewed as hypergraphs. Random instances of k-CSPs, which are also highly expanding, are well-known to be hard to approximate using known algorithmic techniques (and are widely believed to be hard to approximate in polynomial time). However, we show that this is not necessarily the case for instances where the hypergraph is a *high-dimensional expander*.**

We consider the spectral definition of high-dimensional expansion used by Dinur and Kaufman [FOCS 2017] to construct certain primitives related to PCPs. They measure the expansion in terms of a parameter γ which is the analogue of the second singular value for expanding graphs. Extending the results by Barak, Raghavendra and Steurer [FOCS 2011] for 2-CSPs, we show that if an instance of MAX k-CSP over alphabet $[q]$ is a high-dimensional expander with parameter γ, then it is possible to approximate the maximum fraction of satisfiable constraints up to an additive error ε using $q^{O(k)} \cdot (k/\varepsilon)^{O(1)}$ levels of the sum-of-squares SDP hierarchy, provided $\gamma \le \varepsilon^{O(1)} \cdot (1/(kq))^{O(k)}$.

Based on our analysis, we also suggest a notion of threshold-rank for hypergraphs, which can be used to extend the results for approximating 2-CSPs on low threshold-rank graphs. We show that if an instance of MAX k-CSP has threshold rank r for a threshold $\tau = (\varepsilon/k)^{O(1)} \cdot (1/q)^{O(k)}$, then it is possible to approximately solve the instance up to additive error ε, using $r \cdot q^{O(k)} \cdot (k/\varepsilon)^{O(1)}$ levels of the sum-of-squares hierarchy. As in the case of graphs, high-dimensional expanders (with sufficiently small γ) have threshold rank 1 according to our definition.

***Keywords*-CSP; HDX; SOS;**

I. Introduction

We consider the problem of approximately solving constraint satisfaction problems (CSPs) on instances satisfying certain expansion properties. The role of expansion in understanding the approximability of CSPs with two variables in each constraint (2-CSPs) has been extensively studied and has led to several results, which can also be viewed as no-go results for PCP constructions (since PCPs are hard instances of CSPs). It was shown by Arora et al. [AKK+08] (and strengthened by Makarychev and Makarychev [MM11]) that the Unique Games problem is easily approximable on expanding instances, thus proving that the Unique Games Conjecture of Khot [Kho02] cannot be true for expanding instances. Their results were extended to all 2-CSPs and several partitioning problems in works by Barak, Raghavendra and Steurer [BRS11], Guruswami and Sinop [GS11], and Oveis Gharan and Trevisan [OGT15] under much weaker notions of expansion.

We consider the following question:

When are expanding instances of k-CSPs easy for $k > 2$?

At first glance, the question does not make much sense, since random instances of k-CSPs (which are also highly expanding) are known to be hard for various models of computation (see [KMOW17] for an excellent survey). However, while the kind of expansion exhibited by random instances of CSPs is useful for constructing codes, it is not sufficient for constructing primitives for PCPs, such as locally testable codes [BSHR05]. On the other hand, objects such as high-dimensional expanders, which possess a form of "structured multi-scale expansion" have been useful in constructing derandomized direct-product and direct-sum tests (which can be viewed as locally testable distance amplification codes) [DK17], lattices with large distance [KM18], list-decodable direct product codes [DHK+18], and are thought to be intimately connected with PCPs [DK17]. Thus, from the PCP perspective, it is more relevant to ask if this form of expansion can be used to efficiently approximate constraint satisfaction problems.

Connections to coding theory. Algorithmic results related to expanding CSPs are also relevant for the problem of *decoding* locally testable codes. Consider a code C constructed via k-local operations (such as k-fold direct-sum) on a base code C_0 with smaller distance. Then, a codeword in C is simply an instance of a CSP, where each bit places a constraint on k bits (which is k-XOR in case of direct sum) of the relevant codeword in C_0. The task of decoding a noisy codeword is then equivalent to finding an assignment in C_0, satisfying the maximum number of constraints for the above instance. Thus, algorithms for solving CSPs on expanding instances may lead to new decoding algorithms for codes obtained by applying local operations to a base code. In fact, the list decoding algorithm for direct-product codes by Dinur et al. [DHK+18] also relied on algorithmic results for expanding unique games. Since all constructions of locally testable codes need to have at least some weak expansion [DK12], it is interesting to understand what notions of expansion are amenable to algorithmic techniques.

High-dimensional expanders and our results. A d-dimensional expander is a downward-closed hypergraph (simplicial complex), say X, with edges of size at most $d+1$, such that for every hyperedge $\mathfrak{a} \in X$ (with $|\mathfrak{a}| \leq d-1$), a certain "neighborhood graph" $G(X_{\mathfrak{a}})$ is a spectral expander[1]. Here, the graph $G(X_{\mathfrak{a}})$ is defined to have the vertex set $\{i \mid \mathfrak{a} \cup \{i\} \in X\}$ and edge-set $\{i, j \mid \mathfrak{a} \cup \{i, j\} \in X\}$. If the (normalized) second singular value of each of the neighborhood graph is bounded by γ, X is said to be a γ-high-dimensional expander (γ-HDX).

Note that (the downward closure of) a random sparse $(d+1)$-uniform hypergraph, say with n vertices and $c \cdot n$ edges, is very unlikely to be a d-dimensional expander. With high probability, no two hyperedges share more than one vertex and thus for any $i \in [n]$, the neighborhood graph G_i is simply a disjoint union of cliques of size d, which is very far from an expander. While random hypergraphs do not yield high-dimensional expanders, such objects are indeed known to exists via (surprising) algebraic constructions [LSV05b], [LSV05a], [KO18a], [CTZ18] and are known to have several interesting properties and applications [KKL16], [DHK+18], [KM17], [KO18b], [DDFH18], [DK17], [PRT16].

Expander graphs can simply be thought of as the one-dimensional case of the above definition. The results of Barak, Raghavendra and Steurer [BRS11] for 2-CSPs yield that if the constraint graph of a 2-CSP instance (with size n and alphabet size q) is a sufficiently good (one dimensional) spectral expander, then one can efficiently find solutions satisfying $\mathsf{OPT} - \varepsilon$ fraction of constraints, where OPT denotes the maximum fraction of constraints satisfiable by any assignment. Their algorithm is based on $(q/\varepsilon)^{O(1)}$ levels of the Sum-of-Squares (SoS) SDP hierarchy, and the expansion requirement on the constraint graph is that the (normalized) second singular value should be at most $(\varepsilon/q)^{O(1)}$. We show a similar result for k-CSPs when the corresponding simplicial complex $X_{\mathfrak{I}}$, which is obtained by including one hyperedge for each constraint and taking a downward closure, is a sufficiently good $(k-1)$-dimensional expander.

[1] While there are several definitions of high-dimensional expanders, we consider the one by Dinur and Kaufman [DK17], which is most closely related to spectral expansion, and was also the one shown to be related to PCP applications. Our results also work for a weaker but more technical definition by Dikstein et al. [DDFH18], which we defer till later.

Theorem I.1 (Informal)**.** *Let $\mathfrak{I}$ be an instance of MAX k-CSP on n variables taking values over an alphabet of size q, and let $\varepsilon > 0$. Let the simplicial complex $X_{\mathfrak{I}}$ be a γ-HDX with $\gamma = \varepsilon^{O(1)} \cdot (1/(kq))^{O(k)}$. Then, there is an algorithm based on $(k/\varepsilon)^{O(1)} \cdot q^{O(k)}$ levels of the Sum-of-Squares hierarchy, which produces an assignment satisfying $\mathsf{OPT} - \varepsilon$ fraction of the constraints.*

Remark I.2. *While the level-t relaxation for MAX k-CSP can be solved in time $(nq)^{O(t)}$ [RW17], the rounding algorithms used by [BRS11] and our work do not need the full power of this relaxation. Instead, they are captured by the "local rounding" framework of Guruswami and Sinop [GS12] who show how to implement a local rounding algorithm based on t levels of the SoS hierarchy, in time $q^{O(t)} \cdot n^{O(1)}$ (where q denotes the alphabet size).*

A complete version of our results with detailed proofs is given in [AJT19].

Our techniques. We start by using essentially the same argument for analyzing the SoS hierarchy as was used by [BRS11] (specialized to the case of expanders). They viewed the SoS solution as giving a joint distribution on each pair of variables forming a constraint, and proved that for sufficiently expanding graphs, these distributions can be made close to product distributions, by conditioning on a small number of variables (which governs the number of levels required). Similarly, we consider the conditions under which joint distributions on k-tuples corresponding to constraints can be made close to product distributions. Since the [BRS11] argument shows how to split a joint distribution into two marginals, we can use it to recursively split a set of size k into two smaller ones (one can think of all splitting operations as forming a binary tree with k leaves).

However, our arguments differ in the kind of expansion required to perform the above splitting operations. In the case of 2-CSPs, one splits along the edges of the constraint graph, and thus we only need the expansion of the contraint graph (which is part of the assumption). However, in the case of k-CSPs, we may split a set of size $(\ell_1 + \ell_2)$ into disjoint sets of size ℓ_1 and ℓ_2. This requires understanding the expansion of the following family of (weighted) bipartite graphs arising from the complex $X_{\mathfrak{I}}$: The vertices in the graph are sets of variables of size ℓ_1 and ℓ_2 that occur in some constraint, and the weight of an edge $\{\mathfrak{a}_1, \mathfrak{a}_2\}$ for $\mathfrak{a}_1 \cap \mathfrak{a}_2 = \emptyset$, is proportional to the probability that a random constraint contains $\mathfrak{a}_1 \sqcup \mathfrak{a}_2$. Note that this graph may be weighted even if the k-CSP instance $\mathfrak{I}$ is unweighted.

We view the above graphs as random walks, which we call "swap walks" on the hyperedges (faces) in the complex X. While several random walks on high-dimensional expanders have been shown to have rapid mixing [KM17], [KO18b], [DK17], [LLP17], we need a stronger condition. To apply the argument from [BRS11], we not only need that the second singular value is bounded away from one, but require it to be an arbitrarily small constant (as a function of ε, k and q). We show that this is indeed ensured by the condition that $\mathfrak{a}_1 \cap \mathfrak{a}_2 = \emptyset$, and obtain a bound of $k^{O(k)} \cdot \gamma$ on the second singular value. This bound, which constitutes much of the technical work in the paper, is obtained by first expressing these walks in terms of more canonical walks, and then using the beautiful machinery of harmonic analysis on expanding posets by Dikstein et al. [DDFH18] to understand their spectra.

The swap walks analyzed above represent natural random walks on simplicial complexes, and their properties may be of independent interest for other applications. Just as the high-dimensional expanders are viewed as "derandomized" versions of the complete complex (containing all sets of size at most k), one can view the swap walks as derandomized versions of (bipartite) Kneser graphs, which have vertex sets $\binom{[n]}{\ell_1}$ and $\binom{[n]}{\ell_2}$, and edges $(\mathfrak{a}, \mathfrak{b})$ iff $\mathfrak{a} \cap \mathfrak{b} = \emptyset$. We provide a more detailed and technical overview in Section III after discussing the relevant preliminaries in Section II.

High-dimensional threshold rank. The correlation breaking method in [BRS11] can be applied as long as the graph has low threshold rank i.e., the number of singular values above a threshold $\tau = (\varepsilon/q)^{O(1)}$ is bounded. Similarly, the analysis described above can be applied, as long as all the swap walks which arise when splitting the k-tuples have bounded threshold rank. This suggests a notion of high-dimensional threshold rank for hypergraphs (discussed in Section VII), which can be defined in terms of the threshold ranks of the relevant swap walks. We remark that it is easy to show that dense hypergraphs (with $\Omega(n^k)$ hyperedges) have small-threshold rank according to this notion, and thus it can be used to recover known algorithms for approximating k-CSPs on dense instances [FK96] (as was true for threshold rank in graphs).

Other related work. While we extend the approach taken by [BRS11] for 2-CSPs, somewhat different approaches were considered by Guruswami and Sinop [GS11], and Oveis-Gharan and Trevisan [OGT15]. The work by Guruswami and Sinop relied on the expansion of the label extended graph, and used an analysis based on low-dimensional approximations of the SDP solution. Oveis-Gharan and Trevisan used low-threshold rank assumptions to obtain a regularity lemma, which was then used to approximate the CSP. For the case of k-CSPs, the Sherali-Adams hierarchy can be used to solve instances with bounded treewidth [WJ04] and approximately dense instances [YZ14], [MR17]. Brandao and Harrow [BH13] also extended the results by [BRS11] for 2-CSPs to the case of 2-local Hamiltonians. We show that their ideas can also be used to prove a similar extension of our results to k-local Hamiltonians on high-dimensional expanders.

In case of high-dimensional expanders, in addition to canonical walks described here, a "non-lazy" version of these walks (moving from $\mathfrak{s}$ to $\mathfrak{t}$ only if $\mathfrak{s} \neq \mathfrak{t}$) was also considered by Kaufman and Oppenheim [KO18b], Anari et al. [ALGV18] and Dikstein et al. [DDFH18]. The swap walks studied in this paper were also considered independently in a very recent work of Dikstein and Dinur [DD19] (under the name "complement walks").

In a recent follow-up work [AJQ+19], the algorithms developed here were also used to obtain new unique and list decoding algorithms for direct sum and direct product codes, obtained by "lifting" a base code C_0 via k-local operations to amplify distance. This work also showed that the hypergraphs obtained by considering collections of length-k walks on an expanding graph also satisfy (a slight variant of) splittability, and admit similar algorithms.

II. Preliminaries and Notation

A. Linear Algebra

Recall that for an operator $\mathsf{A} : V \to W$ between two finite-dimensional inner product spaces V and W, the operator norm can be written as

$$\|\mathsf{A}\|_{\text{op}} = \sup_{f,g \neq 0} \frac{\langle \mathsf{A}f, g \rangle}{\|f\| \, \|g\|} .$$

Also, for such an A the adjoint $\mathsf{A}^\dagger : W \to V$ is defined as the (unique) operator satisfying $\langle \mathsf{A}f, g \rangle = \langle f, \mathsf{A}^\dagger g \rangle$ for all $f \in V, g \in W$. For $\mathsf{A} : V \to W$, we take $\|\mathsf{A}\|_{\text{op}} = \sigma_1(\mathsf{A}) \geq \sigma_2(\mathsf{A}) \geq \cdots \geq \sigma_r(\mathsf{A}) > 0$ to be its singular values in descending order. Note that for $\mathsf{A} : V \to V$, $\sigma_2(\mathsf{A})$ denotes its second largest eigenvalue in absolute value.

B. High-Dimensional Expanders

A high-dimensional expander (HDX) is a particular kind of downward-closed hypergraph (simplicial complex) satisfying an expansion requirement. We elaborate on these properties and define well known natural walks on HDXs below.

1) Simplicial Complexes:

Definition II.1. *A* simplicial complex *X with ground set $[n]$ is a downward-closed collection of subsets of $[n]$ i.e., for all sets $\mathfrak{s} \in X$ and $\mathfrak{t} \subseteq \mathfrak{s}$, we also have $\mathfrak{t} \in X$. The sets in X are also referred to as* faces *of X.*

We use the notation $X(i)$ to denote the collection of all faces $\mathfrak{s}$ in X with $|\mathfrak{s}| = i$. When faces are of cardinality at most d, we also use the notation $X(\leq d)$ to denote all the faces of X. By convention, we take $X(0) := \{\varnothing\}$.

A simplicial complex $X(\leq d)$ is said to be a pure simplicial complex *if every face of X is contained in some face of size d. Note that in a pure simplicial complex $X(\leq d)$, the top slice $X(d)$ completely determines the complex.*

Note that it is more common to associate a geometric representation to simplicial complexes, with faces of cardinality i being referred to as faces of *dimension* $i - 1$ (and the collection being denoted by $X(i-1)$ instead of

$X(i)$). However, since we will only be treating these as hypergraphs, we prefer to index faces by their cardinality, to improve readability of related expressions.

An important simplicial complex is the complete complex.

Definition II.2 (Complete Complex $\Delta_d(n)$). *We denote by $\Delta_d(n)$ the complete complex with faces of size at most d i.e., $\Delta_d(n) \;:=\; \{\mathfrak{s} \subseteq [n] \mid |\mathfrak{s}| \le d\}$.*

2) Walks and Measures on Simplicial Complexes: Let C^k denote the space of real valued functions on $X(k)$ i.e.,

$$C^k \;:=\; \{f \mid f\colon X(k) \to \mathbb{R}\} \cong \mathbb{R}^{X(k)}.$$

We describe natural walks on simplicial complexes considered in [DK17], [DDFH18], [KO18b], as stochastic operators, which map functions in C^i to C^{i+1} and vice-versa.

To define the stochastic operators associated with the walks, we first need to describe a set of probability measures which serve as the stationary measures for these random walks. For a pure simplicial complex $X(\le d)$, we define a collection of probability measures $(\Pi_1, \ldots \Pi_d)$, with Π_i giving a distribution on faces in the slice $X(i)$.

Definition II.3 (Probability measures $(\Pi_1, \ldots, \Pi_d)$). *Let $X(\le d)$ be a pure simplicial complex and let Π_d be an arbitrary probability measure on $X(d)$. We define a coupled array of random variables $(\mathfrak{s}^{(d)}, \ldots, \mathfrak{s}^{(1)})$ as follows: sample $\mathfrak{s}^{(d)} \sim \Pi_d$ and (recursively) for each $i \in [d]$, take $\mathfrak{s}^{(i-1)}$ to be a uniformly random subset of $\mathfrak{s}^{(i)}$, of size $i-1$.*

The distributions $\Pi_{d-1}, \ldots, \Pi_1$ are then defined to be the marginal distributions of the random variables $\mathfrak{s}^{(d-1)}, \ldots, \mathfrak{s}^{(1)}$ as defined above.

The following is immediate from the definition above.

Proposition II.4. *Let $\mathfrak{a} \in X(\ell)$ be an arbitrary face. For all $j \ge 0$, one has*

$$\sum_{\substack{\mathfrak{b} \in X(\ell+j):\\ \mathfrak{b} \supseteq \mathfrak{a}}} \Pi_{\ell+j}(\mathfrak{b}) \;=\; \binom{\ell+j}{j} \cdot \Pi_\ell(\mathfrak{a}).$$

For all k, we define the inner product of functions $f, g \in C^k$, according to associated measure Π_k

$$\langle f, g\rangle = \mathop{\mathbb{E}}_{\mathfrak{s}\sim\Pi_k}[f(\mathfrak{s})g(\mathfrak{s})] = \sum_{\mathfrak{s}\in X(k)} f(\mathfrak{s})g(\mathfrak{s}) \cdot \Pi_k(\mathfrak{s})\,.$$

We now define the up and down operators $\mathsf{U}_i : C^i \to C^{i+1}$ and $\mathsf{D}_{i+1} : C^{i+1} \to C^i$ as

$$\begin{aligned}
[\mathsf{U}_i g](\mathfrak{s}) \;&=\; \mathop{\mathbb{E}}_{\substack{\mathfrak{s}' \in X(i),\\ \mathfrak{s}' \subseteq \mathfrak{s}}}[g(\mathfrak{s}')] \;=\; \frac{1}{i+1} \cdot \sum_{x \in \mathfrak{s}} g(\mathfrak{s}\backslash\{x\}) \\
[\mathsf{D}_{i+1} g](\mathfrak{s}) \;&=\; \mathop{\mathbb{E}}_{\mathfrak{s}' \sim \Pi_{i+1} | \mathfrak{s}' \supset \mathfrak{s}}[g(\mathfrak{s}')]\,, \\
&=\; \frac{1}{i+1} \cdot \sum_{x \notin \mathfrak{s}} g(\mathfrak{s} \sqcup \{x\}) \cdot \frac{\Pi_{i+1}(\mathfrak{s} \sqcup \{x\})}{\Pi_i(\mathfrak{s})}.
\end{aligned}$$

An important consequence of the above definition is that U_i and D_{i+1} are adjoints with respect to the inner products of C^i and C^{i+1}.

Fact II.5. $\mathsf{U}_i = \mathsf{D}_{i+1}^{\dagger}$, *i.e., $\langle \mathsf{U}_i f, g\rangle = \langle f, \mathsf{D}_{i+1} g\rangle$ for every $f \in C^i$ and $g \in C^{i+1}$.*

Note that the operators can be thought of as defining random walks in a simplicial complex $X(\le d)$. The operator U_i moves *down* from a face $\mathfrak{s} \in X(i+1)$ to a face $\mathfrak{s}' \in X(i)$, but lifts a function $g \in C^i$ *up* to a function $\mathsf{U}g \in C^{i+1}$. Similarly, the operator D_{i+1} can be thought of as defining a random walk which moves *up* from $\mathfrak{s} \in X(i)$ to $\mathfrak{s}' \in X(i+1)$. It is easy to verify that these walks respectively map the measure Π_{i+1} to Π_i, and Π_i to Π_{i+1}.

3) High-Dimensional Expansion: We recall the notion of high-dimensional expansion (defined via local spectral expansion) considered by [DK17]. We first need a few pieces of notation.

For a complex $X(\le d)$ and $\mathfrak{s} \in X(i)$ for some $i \in [d]$, we denote by $X_\mathfrak{s}$ the link complex

$$X_\mathfrak{s} \;:=\; \{\mathfrak{t}\backslash\mathfrak{s} \mid \mathfrak{s} \subseteq \mathfrak{t} \in X\}\,.$$

When $|\mathfrak{s}| \le d-2$, we also associate a natural weighted graph $G(X_\mathfrak{s})$ to a link $X_\mathfrak{s}$, with vertex set $X_\mathfrak{s}(1)$ and edge-set $X_\mathfrak{s}(2)$. The edge-weights are taken to be proportional to the measure Π_2 on the complex $X_\mathfrak{s}$, which is in turn proportional to the measure $\Pi_{|\mathfrak{s}|+2}$ on X. The graph $G(X_\mathfrak{s})$ is referred to as the skeleton of $X_\mathfrak{s}$. Dinur and Kaufman [DK17] define high-dimensional expansion in terms of spectral expansion of the skeletons of the links.

Definition II.6 (γ-HDX from [DK17]). *A simplicial complex $X(\le d)$ is said to be γ-High Dimensional Expander (γ-HDX) if for every $0 \le i \le d-2$ and for every $\mathfrak{s} \in X(i)$, the graph $G(X_\mathfrak{s})$ satisfies $\sigma_2(G(X_\mathfrak{s})) \le \gamma$, where $\sigma_2(G(X_\mathfrak{s}))$ denotes the second singular value of the (normalized) adjacency matrix of $G(X_\mathfrak{s})$.*

C. Constraint Satisfaction Problems (CSPs)

A k-CSP instance $\mathfrak{I} = (H, \mathcal{C}, w)$ with alphabet size q consists of a k-uniform hypergraph, a set of constraints

$$\mathcal{C} = \{\mathcal{C}_\mathfrak{a} \subseteq [q]^\mathfrak{a} : \mathfrak{a} \in H\},$$

and a non-negative weight function $w \in \mathbb{R}_+^H$ on the constraints, satisfying $\sum_{\mathfrak{a}\in H} w(a) = 1$.

A constraint $\mathcal{C}_\mathfrak{a}$ is said to be satisfied by an assignment η if we have $\eta|_\mathfrak{a} \in \mathcal{C}_\mathfrak{a}$ i.e., the restriction of η on $\mathfrak{a}$ is contained in $\mathcal{C}_\mathfrak{a}$. We write, $\mathsf{SAT}_\mathfrak{I}(\eta)$ for the (weighted fraction of the constraints) satisfied by the assignment η i.e.,

$$\mathsf{SAT}_\mathfrak{I}(\eta) \;=\; \sum_{\mathfrak{a}\in H} w(\mathfrak{a}) \cdot \mathbf{1}[\eta|_\mathfrak{a} \in \mathcal{C}_\mathfrak{a}] \;=\; \mathop{\mathbb{E}}_{\mathfrak{a}\sim w} \mathbf{1}[\eta|_\mathfrak{a} \in \mathcal{C}_\mathfrak{a}]\,.$$

We denote by $\mathsf{OPT}(\mathfrak{I})$ the maximum of $\mathsf{SAT}_\mathfrak{I}(\eta)$ over all $\eta \in [q]^{V(H)}$.

Any k-uniform hypergraph H can be associated with a pure simplicial complex in a canonical way by just setting $X_{\mathfrak{I}} = \{\mathfrak{b} : \exists\, \mathfrak{a} \in H \text{ and } \mathfrak{a} \supseteq \mathfrak{b}\}$ – notice that $X_{\mathfrak{I}}(k) = H$. We will refer to this complex as the constraint complex of the instance $\mathfrak{I}$. The probability distribution Π_k on $X_{\mathfrak{I}}$ will be derived from the weights function w of the constraint, i.e

$$\Pi_k(\mathfrak{a}) = w(\mathfrak{a}) \quad \forall \mathfrak{a} \in X_{\mathfrak{I}}(k) = H.$$

D. Sum-of-Squares Relaxations and t-local PSD Ensembles

The Sum-of-Squares (SoS) hierarchy gives a sequence of increasingly tight semidefinite programming relaxations for several optimization problems, including CSPs. Since we will use relatively few facts about the SoS hierarchy, already developed in the analysis of Barak, Raghavendra and Steurer [BRS11], we will adapt their notation of t-local distributions to describe the relaxations. For a k-CSP instance $\mathfrak{I} = (H, \mathcal{C}, w)$ on n variables, we consider the following semidefinite relaxation given by t-levels of the SoS hierarchy, with vectors $v_{(S,\alpha)}$ for all $S \subseteq [n]$ with $|S| \le t$, and all $\alpha \in [q]^S$. Here, for $\alpha_1 \in [q]^{S_1}$ and $\alpha_2 \in [q]^{S_2}$, $\alpha_1 \circ \alpha_2 \in [q]^{S_1 \cup S_2}$ denotes the partial assignment obtained by concatenating α_1 and α_2.

$$\begin{array}{llll}
\text{maximize} & \mathop{\mathbb{E}}_{\mathfrak{a}\sim w}\left[\sum_{\alpha\in C_{\mathfrak{a}}} \|v_{(\mathfrak{a},\alpha)}\|^2\right] & =: \mathsf{SDP}(\mathfrak{I}) & \\
\text{subject to} & \langle v_{(S_1,\alpha_1)}, v_{(S_2,\alpha_2)}\rangle & = 0 & \forall\ \alpha_1|_{S_1\cap S_2} \neq \alpha_2|_{S_1\cap S_2} \\
& \langle v_{(S_1,\alpha_1)}, v_{(S_2,\alpha_2)}\rangle & = \langle v_{(S_3,\alpha_3)}, v_{(S_4,\alpha_4)}\rangle & \forall\ S_1\cup S_2 = S_3\cup S_4, \\
& & & \alpha_1\circ\alpha_2 = \alpha_3\circ\alpha_4 \\
& \sum_{j\in[q]} \|v_{(\{i\},j)}\|^2 & = 1 & \forall i\in[n] \\
& \|v_{(\emptyset,\emptyset)}\| & = 1 &
\end{array}$$

For any set S with $|S| \le t$, the vectors $v_{(S,\alpha)}$ induce a probability distribution μ_S over $[q]^S$ such that the assignment $\alpha \in [q]^S$ appears with probability $\|v_{(S,\alpha)}\|^2$. Moreover, these distributions are consistent on intersections i.e., for $T \subseteq S \subseteq [n]$, we have $\mu_{S|T} = \mu_T$, where $\mu_{S|T}$ denotes the restriction of the distribution μ_S to the set T. We use these distributions to define a collection of random variables $\mathbf{Y}_1, \ldots, \mathbf{Y}_n$ taking values in $[q]$, such that for any set S with $|S| \le t$, the collection of variables $\{\mathbf{Y}_i\}_{i\in S}$ have a joint distribution μ_S. Note that the entire collection $(\mathbf{Y}_1, \ldots, \mathbf{Y}_n)$ *may not* have a joint distribution: this property is only true for sub-collections of size t. We will refer to the collection $(\mathbf{Y}_1, \ldots, \mathbf{Y}_n)$ as a t-local ensemble of random variables.

We also have that that for any $T \subseteq [n]$ with $|T| \le t-2$, and any $\beta \in [q]^T$, we can define a $(t-|T|)$-local ensemble $(\mathbf{Y}'_1, \ldots, \mathbf{Y}'_n)$ by "conditioning" the local distributions on the event $\mathbf{Y}_T = \beta$, where Y_T is shorthand for the collection $\{\mathbf{Y}_i\}_{i\in T}$. For any S with $|S| \le t - |T|$, we define the distribution of $\mathbf{Y}'_S$ as $\mu'_S := \mu_{S\cup T}|\{\mathbf{Y}_T = \beta\}$. Finally, the semidefinite program also ensures that for any such conditioning, the conditional covariance matrix

$$\mathsf{M}_{(S_1,\alpha_1)(S_2,\alpha_2)} \;=\; \mathrm{Cov}\left(\mathbf{1}[\mathbf{Y}'_{S_1} = \alpha_1], \mathbf{1}[\mathbf{Y}'_{S_2} = \alpha_2]\right)$$

is positive semidefinite, where $|S_1|, |S_2| \le (t-|T|)/2$. Here, for each pair S_1, S_2 the covariance is computed using the joint distribution $\mu'_{S_1\cup S_2}$. The PSD-ness be easily verified by noticing that the above matrix can be written as the Gram matrix of the vectors

$$w_{(S,\alpha)} \;:=\; \frac{1}{\|v_{(T,\beta)}\|}\cdot v_{(T\cup S,\beta\circ\alpha)} \;-\; \frac{\|v_{(T\cup S,\beta\circ\alpha)}\|^2}{\|v_{(T,\beta)}\|^3}\cdot v_{(T,\beta)}$$

In this paper, we will only consider t-local ensembles such that for every conditioning on a set of size at most $t-2$, the conditional covariance matrix is PSD. We will refer to these as t-local PSD ensembles. We will also need a simple corollary of the above definitions.

Fact II.7. *Let $(\mathbf{Y}_1, \ldots, \mathbf{Y}_n)$ be a t-local PSD ensemble, and let X be any simplicial complex with $X(1) = [n]$. Then, for all $s \le t/2$, the collection $\{\mathbf{Y}_{\mathfrak{a}}\}_{\mathfrak{a}\in X(\le s)}$ is a (t/s)-local PSD ensemble, where $X(\le s) = \bigcup_{i=1}^{s} X(i)$.*

For random variables $\mathbf{Y}_S$ in a t-local PSD ensemble, we use the notation $\{\mathbf{Y}_S\}$ to denote the distribution of $\mathbf{Y}_S$ (which exists when $|S| \le t$). We also define $\mathrm{Var}[\mathbf{Y}_S]$ as $\sum_{\alpha\in[q]^S} \mathrm{Var}[\mathbf{1}[\mathbf{Y}_S = \alpha]]$.

III. Proof Overview: Approximating MAX 4-XOR

We consider a simple example of a specific k-CSP, which captures most of the key ideas in our proof. Let $\mathfrak{I}$ be an unweighted instance of 4-XOR on n Boolean variables. Let H be a 4-uniform hypergraph on vertex set $[n]$, with a hyperedge corresponding to each constraint i.e., each $\mathfrak{a} = \{i_1, i_2, i_3, i_4\} \in H$ corresponds to a constraint in $\mathfrak{I}$ of the form

$$x_{i_1} + x_{i_2} + x_{i_3} + x_{i_4} \;=\; b_{\mathfrak{a}} \pmod 2,$$

for some $b_{\mathfrak{a}} \in \{0,1\}$. Let X denote the constraint complex for the instance $\mathfrak{I}$ such that $X(1) = [n]$, $X(4) = H$ and let $\Pi_1, \ldots, \Pi_4$ be the associated distributions (with Π_4 being uniform on H).

Local vs global correlation: the BRS strategy. We first recall the strategy used by [BRS11], which also suggests a natural first step for our proof. Given a 2-CSP instance with an associated graph G, and a t-local PSD ensemble $\mathbf{Y}_1, \ldots, \mathbf{Y}_n$ obtained from the SoS relaxation, they consider if the "local correlation" of the ensemble is small across the edges of G (which correspond to constraints) i.e.,

$$\mathop{\mathbb{E}}_{\{i,j\}\sim G} \|\{\mathbf{Y}_i\mathbf{Y}_j\} - \{\mathbf{Y}_i\}\{\mathbf{Y}_j\}\|_1 \;\le\; \varepsilon\,.$$

If the local correlation is indeed small, we easily produce an assignment achieving a value $\mathsf{SDP} - \varepsilon$ in expectation, simply by rounding each variable x_i independently according to the distribution $\{\mathbf{Y}_i\}$. On the other hand, if this is not satisfied, they show (as a special case of their proof) that if G is an expander with second eigenvalue $\lambda \leq c \cdot (\varepsilon^2/q^2)$, then variables also have a high "global correlation", between a typical pair $(i,j) \in [n]^2$. Here, q is the alphabet size and c is a fixed constant. They use this to show that for $(\mathbf{Y}'_1, \ldots, \mathbf{Y}'_n)$ obtained by conditioning on the value of a randomly chosen $\mathbf{Y}_{i_0}$, we have

$$\mathop{\mathbb{E}}_{i}\left[\mathrm{Var}\left[\mathbf{Y}_i\right]\right] - \mathop{\mathbb{E}}_{i_0, \mathbf{Y}_{i_0}} \mathop{\mathbb{E}}_{i}\left[\mathrm{Var}\left[\mathbf{Y}'_i\right]\right] \;\geq\; \Omega(\varepsilon^2/q^2)\,,$$

where the expectations over i and i_0 are both according to the stationary distribution on the vertices of G. Since the variance is bounded between 0 and 1, this essentially shows that the local correlation must be at most ε after conditioning on a set of size $O(q^2/\varepsilon^2)$ (although the actual argument requires a bit more care and needs to condition on a somewhat larger set).

Extension to 4-XOR. As in [BRS11], we check if the t-local PSD ensemble $(\mathbf{Y}_1, \ldots, \mathbf{Y}_n)$ obtained from the SDP solution satisfies

$$\mathop{\mathbb{E}}_{\{i_1,i_2,i_3,i_4\} \in H} \left\|\{\mathbf{Y}_{i_1}\mathbf{Y}_{i_2}\mathbf{Y}_{i_3}\mathbf{Y}_{i_4}\} - \{\mathbf{Y}_{i_1}\}\{\mathbf{Y}_{i_2}\}\{\mathbf{Y}_{i_3}\}\{\mathbf{Y}_{i_4}\}\right\|_1 \leq \varepsilon.$$

As before, independently sampling each x_i from $\{\mathbf{Y}_i\}$ gives an expected value at least $\mathsf{SDP} - \varepsilon$ in this case. If the above inequality is not satisfied, an application of triangle inequality gives

$$\mathop{\mathbb{E}}_{\{i_1,i_2,i_3,i_4\} \in H} \left[\begin{array}{l} \left\|\{\mathbf{Y}_{i_1}\mathbf{Y}_{i_2}\mathbf{Y}_{i_3}\mathbf{Y}_{i_4}\} - \{\mathbf{Y}_{i_1}\mathbf{Y}_{i_2}\}\{\mathbf{Y}_{i_3}\mathbf{Y}_{i_4}\}\right\|_1 \;+ \\ \left\|\{\mathbf{Y}_{i_1}\mathbf{Y}_{i_2}\} - \{\mathbf{Y}_{i_1}\}\{\mathbf{Y}_{i_2}\}\right\|_1 \;+ \\ \left\|\{\mathbf{Y}_{i_3}\mathbf{Y}_{i_4}\} - \{\mathbf{Y}_{i_3}\}\{\mathbf{Y}_{i_4}\}\right\|_1 \end{array}\right] > \varepsilon$$

Symmetrizing over all orderings of $\{i_1, i_2, i_3, i_4\}$, we can write the above as

$$\varepsilon_2 + 2 \cdot \varepsilon_1 \;>\; \varepsilon\,,$$

which gives $\max\{\varepsilon_1, \varepsilon_2\} \geq \varepsilon/3$. Here,

$$\begin{aligned} \varepsilon_1 &:= \mathop{\mathbb{E}}_{\{i_1,i_2\} \sim \Pi_2} \left\|\{\mathbf{Y}_{i_1}\mathbf{Y}_{i_2}\} - \{\mathbf{Y}_{i_1}\}\{\mathbf{Y}_{i_2}\}\right\|_1, \quad \text{and} \\ \varepsilon_2 &:= \mathop{\mathbb{E}}_{\{i_1,i_2,i_3,i_4\} \sim \Pi_4} \left\|\{\mathbf{Y}_{i_1}\mathbf{Y}_{i_2}\mathbf{Y}_{i_3}\mathbf{Y}_{i_4}\} - \{\mathbf{Y}_{i_1}\mathbf{Y}_{i_2}\}\{\mathbf{Y}_{i_3}\mathbf{Y}_{i_4}\}\right\|_1 \\ &= \mathop{\mathbb{E}}_{\{i_1,i_2,i_3,i_4\} \sim \Pi_4} \left\|\left\{\mathbf{Y}_{\{i_1,i_2\}}\mathbf{Y}_{\{i_3,i_4\}}\right\} - \left\{\mathbf{Y}_{\{i_1,i_2\}}\right\}\left\{\mathbf{Y}_{\{i_3,i_4\}}\right\}\right\|_1. \end{aligned}$$

As before, ε_1 measures the local correlation across edges of a weighted graph G_1 with vertex set $X(1) = [n]$ and edge-weights given by Π_2. Also, ε_2 measures the analogous quantity for a graph G_2 with vertex set $X(2)$ (pairs of variables occurring in constraints) and edge-weights given by Π_4.

Recall that the result from [BRS11] can be applied to *any* graph G over variables in a 2-local PSD ensemble, as long as the $\sigma_2(G)$ is small. Since $\{\mathbf{Y}_i\}_{i \in [n]}$ and $\{\mathbf{Y}_{\mathfrak{s}}\}_{\mathfrak{s} \in X(2)}$ are both $(t/2)$-local PSD ensembles (by Fact II.7), we will apply the result to the graph G_1 on the first ensemble and G_2 on the second ensemble. We consider the potential

$$\Phi(\mathbf{Y}_1, \ldots, \mathbf{Y}_n) \;:=\; \mathop{\mathbb{E}}_{i \sim \Pi_1} \mathrm{Var}\left[\mathbf{Y}_i\right] \;+\; \mathop{\mathbb{E}}_{\mathfrak{s} \sim \Pi_2} \mathrm{Var}\left[\mathbf{Y}_{\mathfrak{s}}\right].$$

Since local correlation is large along at least one of the graphs G_1 and G_2, using the above arguments (and the non-decreasing nature of variance under conditioning) it is easy to show that in expectation over the choice of $\{i_0, j_0\} \sim \Pi_2$ and $\beta \in [q]^2$ chosen from $\left\{\mathbf{Y}_{\{i_0,j_0\}}\right\}$, the conditional ensemble $(\mathbf{Y}'_1, \ldots, \mathbf{Y}'_n)$ satisfies

$$\Phi(\mathbf{Y}_1, \ldots, \mathbf{Y}_n) - \mathop{\mathbb{E}}_{i_0, j_0, \beta} \Phi(\mathbf{Y}'_1, \ldots, \mathbf{Y}'_n) \;=\; \Omega(\varepsilon^2)\,,$$

provided G_1 and G_2 satisfy $\sigma_2(G_1), \sigma_2(G_2) \;\leq\; c \cdot \varepsilon^2$ for an appropriate constant c.

The bound on the eigenvalue of G_1 follows simply from the fact that it is the skeleton of X, which is a γ-HDX. Obtaining bounds on the eigenvalues of G_2 and similar higher-order graphs, constitutes much of the technical part of this paper. Note that for a random sparse instance of MAX 4-XOR, the graph G_2 will be a matching with high probability (since $\{i_1, i_2\}$ in a constraint will only be connected to $\{i_3, i_4\}$ in the same constraint). However, we show that in case of a γ-HDX, this graph has second eigenvalue $O(\gamma)$. We analyze these graphs in terms of modified high-dimensional random walks, which we call "swap walks".

We remark that our potential and choice of a "seed set" of variables to condition on, is slightly different from [BRS11]. To decrease the potential function above, we need that for each level $X(i)$ ($i = 1, 2$ in the example above) the seed set must contain sufficiently many independent samples from $X(i)$ sampled according to Π_i. This can be ensured by drawing independent samples from the top level $X(k)$ (though $X(2)$ suffices in the above example). In contrast, the seed set in [BRS11] consists of random samples from Π_1. The graph G_2 defined above can be thought of as a random walk on $X(2)$, which starts at a face $\mathfrak{s} \in X(2)$, moves up to a face (constraint) $\mathfrak{s}' \in X(4)$ containing it, and then descends to a face $\mathfrak{t} \in X(2)$ such that $\mathfrak{t} \subset \mathfrak{s}'$ and $\mathfrak{s} \cap \mathfrak{t} = \emptyset$ i.e., the walk "swaps out" the elements in $\mathfrak{s}$ for other elements in $\mathfrak{s}'$. Several walks considered on simplicial complexes allow for the possibility of a non-trivial intersection, and hence have second eigenvalue lower bounded by a constant. On the other hand, swap walks completely avoid any laziness and thus turn out to have eigenvalues which can be made arbitrarily small. To understand the eigenvalues for this walk, we will express it in terms of other canonical walks defined on simplicial complexes.

Recall that the up and down operators can be used to define random walks on simplicial complexes. The up operator $\mathsf{U}_i : C^i \to C^{i+1}$ defines a walk that moves *down* from a face $\mathfrak{s} \in X(i+1)$ to a random face $\mathfrak{t} \in X(i), \mathfrak{t} \subset \mathfrak{s}$ (the operator thus "lifts" a function in C^i to a function in C^{i+1}). Similarly, the down operator $\mathsf{D}_i : C^i \to C^{i-1}$ moves *up* from a face $\mathfrak{s} \in X(i-1)$ to $\mathfrak{t} \in X(i), \mathfrak{t} \supset s$, with probability $\Pi_i(\mathfrak{t})/(i \cdot \Pi_{i-1}(\mathfrak{s}))$. These can be used to define a canonical random walk

$$\mathsf{N}_{2,2}^{(u)} \;:=\; \mathsf{D}_3 \cdots \mathsf{D}_{u+2}\mathsf{U}_{u+1} \cdots \mathsf{U}_2\,, \qquad \mathsf{N}_{2,2}^{(u)} : C^2 \to C^2\,,$$

which moves from up for u steps $\mathfrak{s} \in X(2)$ to $\mathfrak{s}' \in X(u+2)$, and then descends back to $\mathfrak{t} \in X(2)$. Such walks were analyzed optimally by Dinur and Kaufman [DK17], who proved that $\lambda_2\left(\mathsf{N}_{2,2}^{(u)}\right) = 2/(u+2) \pm O_u(\gamma)$ when X is a γ-HDX. Thus, while this walk gives an expanding graph with vertex set $X(2)$, the second eigenvalue cannot be made arbitrarily small for a fixed u (recall that we are interested in showing that $\sigma_2(G_2) \leq c \cdot \varepsilon^2$). However, note that we are only interested in $\mathsf{N}_{2,2}^{(2)}$ *conditioned on the event* that the two elements from $\mathfrak{s}$ are "swapped out" with new elements in the final set $\mathfrak{t}$ i.e., $\mathfrak{s} \cap \mathfrak{t} = \emptyset$. We define

$$\mathsf{S}_{2,2}^{(u,j)}(\mathfrak{s},\mathfrak{t}) \;:=\; \begin{cases} \frac{\binom{u+2}{2}}{\binom{u}{j}\cdot\binom{2}{2-j}} \cdot \mathsf{N}_{2,2}^{(u)} & \text{if } |\mathfrak{t} \setminus \mathfrak{s}| = j \\ 0 & \text{otherwise} \end{cases},$$

where the normalization is to ensure stochasticity of the matrix. In this notation, the graph G_2 corresponds to the random-walk matrix $\mathsf{S}_{2,2}^{(2,2)}$. We show that while $\sigma_2(\mathsf{N}_{2,2}^{(2)}) \approx 1/2$, we have that $\sigma_2(\mathsf{S}_{2,2}^{(2,2)}) = O(\gamma)$. We first write the canonical walks in terms of the swap walks. Note that

$$\mathsf{N}_{2,2}^{(2)} \;=\; \frac{1}{6}\cdot \mathsf{I} \;+\; \frac{2}{3}\cdot \mathsf{S}_{2,2}^{(2,1)} \;+\; \frac{1}{6}\cdot \mathsf{S}_{2,2}^{(2,2)},$$

since the "descent" step from $\mathfrak{s}' \in X(4)$ containing $\mathfrak{s} \in X(2)$, produces a $\mathfrak{t} \in X(2)$ which "swaps out" $0, 1$ and 2 elements with probabilities $1/6, 2/3$ and $1/6$ respectively. Similarly,

$$\mathsf{N}_{2,2}^{(1)} \;=\; \frac{1}{3}\cdot \mathsf{I} \;+\; \frac{2}{3}\cdot \mathsf{S}_{2,2}^{(1,1)}.$$

Finally, we use the fact (proved in Section IV) that while the canonical walks do depend on the "height" u (i.e., $\mathsf{N}_{2,2}^{(u)} \neq \mathsf{N}_{2,2}^{(u')}$) the swap walks (for a fixed number of swaps j) are independent of the height to which they ascend! In particular, we have

$$\mathsf{S}_{2,2}^{(2,1)} \;=\; \mathsf{S}_{2,2}^{(1,1)}.$$

Using these, we can derive an expression for the swap walk $\mathsf{S}_{2,2}^{(2,2)}$ as

$$\mathsf{S}_{2,2}^{(2,2)} \;=\; \mathsf{I} + 6\cdot \mathsf{N}_{2,2}^{(2)} - 6\cdot \mathsf{N}_{2,2}^{(1)} \;=\; \mathsf{I} + 6\cdot(\mathsf{D}_3\mathsf{D}_4\mathsf{U}_3\mathsf{U}_2 - \mathsf{D}_3\mathsf{U}_2)$$

To understand the spectrum of operators such as the ones given by the above expression, we use the beautiful machinery for harmonic analysis over HDXs (and more generally over expanding posets) developed by Dikstein et al. [DDFH18]. They show how to decompose the spaces C^k into approximate eigenfunctions for operators of the form DU. Using these decompositions and the properties of expanding posets, we can show that distinct eigenvalues of the above operator are approximately the same (up to $O(\gamma)$ errors) when analyzing the walks on the complete complex. Finally, we use the fact that swap walks in a complete complex correspond to Kneser graphs (for which the eigenvectors and eigenvalues are well-known) to show that $\lambda_2(\mathsf{S}_{2,2}^{(2,2)}) = O(\gamma)$.

Splittable CSPs and high-dimensional threshold rank. We note that the ideas used above can be generalized (at least) in two ways. In the analysis of distance from product distribution for a 4-tuple of random variables forming a contraint, we split it in 2-tuples. In general, we can choose to split tuples in a k-CSP instance along *any* binary tree $\mathcal{T}$ with k leaves, with each parent node corresponding to a swap walk between tuples forming its children. Finally, the analysis from [BRS11] also works if the each of the swap walks in some $\mathcal{T}$ have a bounded number (say r) of eigenvalues above some threshold τ, which provide a notion of high-dimensional threshold rank for hypergraphs. We refer to such an instance as a $(\mathcal{T}, \tau, r)$-splittable.

The arguments sketched above show that high-dimensional expanders are $(\mathcal{T}, O(\gamma), 1)$-splittable for all $\mathcal{T}$. Since the knowledge of $\mathcal{T}$ is only required in our analysis and not in the algorithm, we say that $\mathrm{rank}_\tau(\mathfrak{I}) \leq r$ (or that $\mathfrak{I}$ is (τ, r)-splittable) if $\mathfrak{I}$ is $(\mathcal{T}, \tau, r)$-splittable for any $\mathcal{T}$. We defer the precise statement of results for (τ, r)-splittable instances to Section VII.

IV. Walks

It is important to note that both U_i and D_{i+1} are row-stochastic matrices, i.e. we can think of them as the probability matrices describing the movement of a walk from $X(i+1)$ to $X(i)$; and from $X(i)$ to $X(i+1)$ respectively.

Concretely, we will think $[\mathsf{D}_{i+1}^\top e_{\mathfrak{s}}](\mathfrak{t})$ as the probabliity of the walk moving up from $\mathfrak{s} \in X(i)$ to $\mathfrak{t} \in X(i+1)$. Similarly, we will think of $[\mathsf{U}_i^\top e_{\mathfrak{t}}](\mathfrak{s})$ as the walk moving down from $\mathfrak{t} \in X(i+1)$ to $\mathfrak{s} \in X(i)$.

By referring to the definition of the up and down operators in Section II, we observe

$$[\mathsf{D}_{i+1}^\top e_{\mathfrak{s}}](\mathfrak{t}) = \frac{\mathbf{1}[\mathfrak{t} \supseteq \mathfrak{s}]}{i+1}\frac{\Pi_{i+1}(\mathfrak{t})}{\Pi_i(\mathfrak{s})} \text{ and } [\mathsf{U}_i^\top e_{\mathfrak{t}}](\mathfrak{s}) = \frac{\mathbf{1}[\mathfrak{s} \subseteq \mathfrak{t}]}{i+1}.$$

With this observation it is easy to see that our notion of random walk respects the probability distributions Π_j, i.e. we have

$$\mathsf{U}_i^\top \Pi_{i+1} = \Pi_i \quad \text{and} \quad \mathsf{D}_{i+1}^\top \Pi_i = \Pi_{i+1},$$

i.e., randomly moving up from a sample of Π_j gives a sample of Π_{j+1} and similarly, moving down from a sample of Π_{j+1} results in a sample of Π_j.

Instead of going up and down by one dimension, one can try going up or down by multiple dimensions since $(\mathsf{D}_{i+1} \cdots \mathsf{D}_{i+\ell})$ and $(\mathsf{U}_{i+\ell} \cdots \mathsf{U}_i)$ are still row-stochastic matrices. Further, the corresponding probability vectors still have intuitive explanations in terms of the distributions Π_j. For a face $\mathfrak{s} \in X(k)$, we introduce the notation

$$p_{\mathfrak{s}}^{(u)} = (\mathsf{D}_{k+1} \cdots \mathsf{D}_{k+u})^\top e_{\mathfrak{s}}$$

where we take $p_{\mathfrak{s}}^{(0)} = e_{\mathfrak{s}}$. This notation will be used to denote the probability distribution of the up-walk starting from $\mathfrak{s} \in X(k)$ and ending in a random face $\mathfrak{t} \in X(k+u)$ satisfying $\mathfrak{t} \supseteq \mathfrak{s}$.

Note that the following Lemma together with Proposition II.4 implies that $p_{\mathfrak{s}}^{(u)}$ is indeed a probability distribution.

Proposition IV.1. *For $\mathfrak{s} \in X(k)$ and $\mathfrak{a} \in X(k+u)$ one has,*

$$p_{\mathfrak{s}}^{(u)}(\mathfrak{a}) = \mathbf{1}[\mathfrak{a} \supseteq \mathfrak{s}] \cdot \frac{1}{\binom{k+u}{u}} \cdot \frac{\Pi_{k+u}(\mathfrak{a})}{\Pi_k(\mathfrak{s})}.$$

Similarly, we introduce the notation $q_{\mathfrak{a}}^{(u)}$, as

$$q_{\mathfrak{a}}^{(u)}(\mathfrak{s}) = (\mathsf{U}_{k+u-1} \cdots \mathsf{U}_k)^\top e_{\mathfrak{s}},$$

i.e. for the probability distribution of the down-walk starting from $\mathfrak{a} \in X(k+u)$ and ending in a random face of $X(k)$ contained in $\mathfrak{a}$. The following can be verified using Proposition IV.1, and the fact that $(\mathsf{U}_{k+u-1} \cdots \mathsf{U}_k)^\dagger = \mathsf{D}_{k+u} \cdots \mathsf{D}_{k+1}$.

Corollary IV.2. *Let $X(\leq d)$ be a simplicial complex, and $k, u \geq 0$ be parameters satisfying $k+u \leq d$. For $\mathfrak{a} \in X(k+u)$ and $\mathfrak{s} \in X(k)$, one has*

$$q_{\mathfrak{a}}^{(u)}(\mathfrak{s}) = \frac{1}{\binom{k+u}{u}} \cdot \mathbf{1}[\mathfrak{s} \subseteq \mathfrak{a}].$$

In the remainder of this section, we will try to construct more intricate walks on X from $X(k)$ to $X(l)$.

A. The Canonical and the Swap Walks on a Simplicial Complex

Definition IV.3 (Canonical and Swap u-Walks)**.** *Let $d \geq 0$, $X(\leq d)$ be a simplicial complex, and $k, l, u \geq 0$ be parameters satisfying $l \leq k$, $u \leq l$ and $d \geq k+u$; where the constraints on these parameters are to ensure well-definedness. We will define the following random walks,*

- ***canonical u-walk from $X(k)$ to $X(l)$**. Let $\mathsf{N}_{k,l}^{(u)}$ be the (row-stochastic) Markov operator that represents the following random walk: Starting from a face $\mathfrak{s} \in X(k)$,*
  - (random ascent/up-walk) *randomly move up a face $\mathfrak{s}'' \in X(k+u)$ that contains $\mathfrak{s}$, where $\mathfrak{s}''$ is picked with probability*

    $$p_{\mathfrak{s}}^{(u)}(\mathfrak{s}'') = [(\mathsf{D}_{k+1} \cdots \mathsf{D}_{k+u})^\top e_{\mathfrak{s}}](\mathfrak{s}'').$$

  - (random descent/down-walk) *go to a face $\mathfrak{s}' \in X(l)$ picked uniformly among all the l-dimensional faces that are contained in $\mathfrak{s}''$, i.e., the set $\mathfrak{s}'$ is picked with probability*

    $$q_{\mathfrak{s}''}(\mathfrak{s}') = \frac{\mathbf{1}[\mathfrak{s}' \subseteq \mathfrak{s}'']}{\binom{k+u}{l}} = [(\mathsf{U}_{k+u-1} \cdots \mathsf{U}_l)^\top e_{\mathfrak{s}''}](\mathfrak{s}').$$

  The operator $\mathsf{N}_{k,l}^{(u)}\colon C^l \to C^k$ satisfies the following equation,

  $$\mathsf{N}_{k,l}^{(u)} = \mathsf{D}_{k+1} \cdots \mathsf{D}_{k+u} \mathsf{U}_{k+u-1} \cdot \mathsf{U}_k \cdots \mathsf{U}_l.$$

  Notice that we have $\mathsf{N}_{k,k}^{(0)} = \mathsf{I}$, and $\mathsf{N}_{k,l}^{(0)} = (\mathsf{U}_{k-1} \ldots \mathsf{U}_l)$ for $l < k$.
- ***swapping walk from $X(k)$ to $X(l)$**. Let $\mathsf{S}_{k,l}$ be the Markov operator that represents the following random walk: Starting from a face $\mathfrak{s} \in X(k)$,*
  - (random ascent/up-walk) *randomly move up to a face $\mathfrak{s}'' \in X(k+l)$ that contains $\mathfrak{s}$, where as before $\mathfrak{s}''$ is picked with probability*

    $$p_{\mathfrak{s}}^{(l)}(\mathfrak{s}'') = [(\mathsf{D}_{k+1} \cdots \mathsf{D}_{k+l+1})^\top e_{\mathfrak{s}}](\mathfrak{s}'').$$

  - (deterministic descent) *deterministically go to $\mathfrak{s}' = \mathfrak{s}'' \backslash \mathfrak{s} \in X(l)$.*

For our applications, we will need to show that the walk $\mathsf{S}_{k,l}$ has good spectral expansion whenever X is a d-dimensional γ-expander, for γ sufficiently small. To show this, we will relate the swapping walk operator $\mathsf{S}_{k,l}$ on X to the canonical random walk operators $\mathsf{N}_{k,l}^{(u)}$ (q.v. Lemma IV.4).

By the machinery of expanding posets (q.v. Section V) it is possible to argue that the spectral expansion of the random walk operator $\mathsf{N}_{k,l}^{(u)}$ on a high dimensional expander will be close to that of the complete complex. This will allow us to conclude using the relation between the swapping walks and the canonical walks (q.v. Lemma IV.4) that the spectral expansion of the swapping walk on X, will be comparable with the spectral expansion of the swap walk on the complete complex. More precisely, we will show

Lemma IV.4 (Lemma V.34)**.** *For any $d, k, l \geq 0$, and the complete simplicial simplicial complex $X(\leq d)$, one has the following: If $k \geq l \geq 0$ and $d \geq k+l$, we have*

$$\sigma_2(\mathsf{S}_{k,l}) = O_{k,l}\left(\frac{1}{n}\right).$$

Using these two, and the expanding poset machinery, we will conclude

Theorem IV.5 (Theorem V.2 simplified)**.** *Let X be a d-dimensional γ expander. If $k \geq l \geq 0$ satisfy $d \geq l+k$ we have,*

$$\sigma_2(\mathsf{S}_{k,l}) = O_{k,l}(\gamma)$$

where $\mathsf{S}_{k,l}$ is the swapping walk on X from $X(k)$ to $X(l)$.

To prove Theorem IV.5 we will need to define an intermediate random walk that we will call the j-swapping u-walk from $X(k)$ to $X(l)$:

Definition IV.6 (j-swapping u-walk from $X(k)$ to $X(l)$)**.** *Given $d, u, j, k, l \geq 0$ satisfying $l \leq k$, $j \leq u$, $u \leq l$, and $d \geq k+u$. Let $\mathsf{S}_{k,l}^{(u,j)}$ be the Markov operator that represents the following random walk from $X(k)$ to $X(l)$ on a d-dimensional simplicial complex X: Starting from $\mathfrak{s} \in X(k)$*

- (random ascent/up-walk) *randomly move up to a face $\mathfrak{s}'' \in X(k+u)$ that contains $\mathfrak{s}$, where $\mathfrak{s}''$ is picked with probability*

$$p^{(u)}_{\mathfrak{s}}(\mathfrak{s}'') = [(\mathsf{D}_{k+1} \cdots \mathsf{D}_{k+u})^\top e_{\mathfrak{s}}](\mathfrak{s}'').$$

- (conditioned descent) *go to a face $\mathfrak{s}' \in X(l)$ sampled uniformly among all the subsets of $\mathfrak{s}'' \in X(k+u)$ that have intersection j with $\mathfrak{s}'' \backslash \mathfrak{s}$, i.e. $|\mathfrak{s}' \cap (\mathfrak{s}'' \backslash \mathfrak{s})| = j$.*

Notice that $\mathsf{S}_{k,l} = \mathsf{S}^{(l,l)}_{k,l}$ for any k and $\mathsf{I} = \mathsf{S}^{(u,0)}_{k,k}$ for any u.

Remark IV.7. *We will prove that the parameter u does not effect the swapping walk $\mathsf{S}^{(u,j)}_{k,l}$ so long as $u \geq j$, i.e. for all $u, u' \geq j$ we have $\mathsf{S}^{(u',j)}_{k,l} = \mathsf{S}^{(u,j)}_{k,l}$. Thus, we will often write $\mathsf{S}^{(j)}_{k,l}$ for $\mathsf{S}^{(j,j)}_{k,l}$.*

B. Swap Walks are Height Independent

Recall that the swap walk $\mathsf{S}^{(u,j)}_{k,l}$ is the conditional walk defined in terms of $\mathsf{N}^{(u)}_{k,l}$ where $\mathfrak{s} \in X(k)$ is connected to $\mathfrak{t} \in X(l)$ only if $|\mathfrak{t} \setminus \mathfrak{s}| = j$. The parameter u is called the height of the walk, namely the number of times it moves up. Since up and down operators have second singular value bounded away from 1, the second singular value of $\mathsf{N}^{(u)}_{k,l}$ shrinks as u increases. In other words, the operator $\mathsf{N}^{(u)}_{k,l}$ depends on the height u. Surprisingly, the walk $\mathsf{S}^{(u,j)}_{k,l}$ which is defined in terms of $\mathsf{N}^{(u)}_{k,l}$ does not depend on the particular choice of u as long as it is well defined. More precisely, we have the following result.

Lemma IV.8. *If X is a d-dimensional simplicial complex, $0 \leq l \leq k$, and $u, u' \in [j, d-k]$, then*

$$\mathsf{S}^{(u,j)}_{k,l} = \mathsf{S}^{(u',j)}_{k,l}.$$

In order to obtain Lemma IV.8, we will need a simple proposition:

Proposition IV.9. *Let $\mathfrak{s} \in X(k)$, $\mathfrak{s}' \subseteq \mathfrak{s}$ and $|\mathfrak{t}'| = j$. Suppose $\mathfrak{s}' \sqcup \mathfrak{t}' \in X(l)$. Then, we have*

$$\mathsf{S}^{(u,j)}_{k,l}(\mathfrak{s}, \mathfrak{s}' \sqcup \mathfrak{t}') = \frac{1}{\binom{k}{l-j} \cdot \binom{u}{j}} \cdot \sum_{\substack{\mathfrak{a} \in X(k+u): \\ \mathfrak{a} \supseteq (\mathfrak{s} \sqcup \mathfrak{t}')}} p^{(u)}_{\mathfrak{s}}(\mathfrak{a}).$$

Lemma IV.10 (Height Independence)**.** *Let $u \in [j, d-k]$. For any $\mathfrak{s} \in X(k)$, $\mathfrak{s}' \subseteq \mathfrak{s}$ and $\mathfrak{t}' \in X(j)$ satisfying $\mathfrak{s}' \sqcup \mathfrak{t}' \in X(l)$ we have the following,*

$$\mathsf{S}^{(u,j)}_{k,l}(\mathfrak{s}, \mathfrak{s}' \sqcup \mathfrak{t}') = \frac{1}{\binom{k}{l-j}\binom{k+j}{j}} \cdot \frac{\Pi_{k+j}(\mathfrak{s} \sqcup \mathfrak{t}')}{\Pi_k(\mathfrak{s})}.$$

In particular, the choice of $u \in [j, d-k]$ does not affect the swap walk.

Since the choice of u does not affect the formula, we obtain Lemma IV.8.

C. Canonical Walks in Terms of the Swap Walks

We show that the canonical walks are given by an average of swap walks with respect to the hypergeometric distribution.

Lemma IV.11. *Let $u, l, k, d \geq 0$ be given satisfying $l \leq k$ and $u \leq l$. Then, we have the following formula for the canonical u-walk on any $X(\leq d)$ satisfying $d \geq k+u$*

$$\mathsf{N}^{(u)}_{k,l} = \sum_{j=0}^{u} \frac{\binom{u}{j}\binom{k}{l-j}}{\binom{k+u}{l}} \cdot \mathsf{S}^{(j)}_{k,l}.$$

D. Inversion: Swap Walks in Terms of Canonical Walks

We show how the swap walks can be obtained as a signed sum of canonical walks. This result follows from binomial inversion which we recall next.

Fact IV.12 (Binomial Inversion, [BS02])**.** *Let $(a_n)_{n \geq 0}, (b_n)_{n \geq 0}$ be arbitrary sequences. Suppose for all $n \geq 0$ we have,*

$$b_n = \sum_{j=0}^{n} \binom{n}{j} \cdot (-1)^j \cdot a_j.$$

Then, we also have

$$a_n = \sum_{j=0}^{n} \binom{n}{j} \cdot (-1)^j \cdot b_j.$$

Corollary IV.13. *Let $k, l, d \geq 0$ be given parameters such that $k + l \leq d$ and $k \geq l$. For any simplicial complex $X(\leq d)$, one has the following formula for the u-swapping walk from $X(k)$ to $X(l)$ in terms of the canonical j-walks:*

$$\binom{k}{l-u}\mathsf{S}^{(u)}_{k,l} = \sum_{j=0}^{u} (-1)^{u-j} \cdot \binom{k+j}{l} \cdot \binom{u}{j} \cdot \mathsf{N}^{(j)}_{k,l}.$$

V. Spectral Analysis of Swap Walks

Swap walks arise naturally in our k-CSPs approximation scheme on HDXs where the running time and the quality of approximation depend on the expansion of these walks. For this reason, we analyze the spectra of swap walks. We show that swap walks $\mathsf{S}_{k,k}$ of γ-HDXs are indeed expanding for γ sufficiently small. More precisely, the first main result of this section is the following.

Theorem V.1 (Swap Walk Spectral Bound)**.** *Let $X(\leq d)$ be a γ-HDX with $d \geq 2k$. Then the second largest singular value $\sigma_2(\mathsf{S}_{k,k})$ of the swap operator $\mathsf{S}_{k,k}$ is*

$$\sigma_2(\mathsf{S}_{k,k}) \;\leq\; \gamma \cdot \left(2^7 \cdot k^4 \cdot 2^{3k} \cdot k^k\right).$$

Theorem V.1 is enough for the analysis of our k-CSP approximation scheme when k is a power of two. However, to analyze general k-CSPs on HDXs we need to understand the spectra of general swap walks $\mathsf{S}_{k,l}$ where k may differ from l. Therefore, we generalize the spectral analysis of $\mathsf{S}_{k,k}$

above to $\mathsf{S}_{k,l}$ obtaining Theorem V.2, our second main result of this section.

Theorem V.2 (Rectangular Swap Walk Spectral Bound)**.** *Suppose $X(\le d)$ is a γ-HDX with $d \ge k+l$ and $k \le l$. Then the largest non-trivial singular value $\sigma_2(\mathsf{S}_{k,l})$ of the swap operator $\mathsf{S}_{k,l}$ is*

$$\sigma_2(\mathsf{S}_{k,l}) \;\le\; \sqrt{\gamma \cdot (2^8 \cdot k^2\ell^2 \cdot 2^{2k+4l} \cdot k^k)}.$$

A. Square Swap Walks $\mathsf{S}_{k,k}$

We prove Theorem V.1 by connecting the spectral structure of $\mathsf{S}_{k,k}$ of general γ-HDXs to the well behaved case of complete simplicial complexes. To distinguish these two cases we denote by $\mathsf{S}^{\Delta}_{k,k}$ the swap $\mathsf{S}_{k,k}$ of complete complexes [2]. In fact, $\mathsf{S}^{\Delta}_{k,k}$ is the random walk operator of the well known Kneser graph $K(n,k)$ (see Definition V.3).

Definition V.3 (Kneser Graph $K(n,k)$ [GM15])**.** *The Kneser graph $K(n,k)$ is the graph $G=(V,E)$ where $V=\binom{[n]}{k}$ and $E=\{\{\mathfrak{s},\mathfrak{t}\} \mid \mathfrak{s}\cap\mathfrak{t}=\emptyset\}$.*

Then at least for complete complexes we know that $\mathsf{S}^{\Delta}_{k,k}$ is expanding. This is a direct consequence of Fact V.4.

Fact V.4 (Kneser Graph [GM15])**.** *The singular values [3] of the Kneser graph $K(n,k)$ are*

$$\binom{n-k-i}{k-i},$$

for $i=0,\ldots,k$.

This means that $\sigma_2(\mathsf{S}^{\Delta}_{k,k}) = O_k(1/n)$ as shown in Claim V.5.

Claim V.5. *Let $d \ge 2k$ and $\Delta_d(n)$ be the complete complex. The second largest singular value $\sigma_2(\mathsf{S}^{\Delta}_{k,k})$ of the swap operator $\mathsf{S}^{\Delta}_{k,k}$ on $\Delta_d(n)$ is*

$$\sigma_2(\mathsf{S}^{\Delta}_{k,k}) \;=\; \frac{k}{n-k},$$

provided $n \ge M_k$ where $M_k \in \mathbb{N}$ only depends on k.

Therefore, if we could claim that $\sigma_2(\mathsf{S}_{k,k})$ of an arbitrary γ-HDX is close to $\sigma_2(\mathsf{S}^{\Delta}_{k,k})$ (provided γ is sufficiently small), we would conclude that general $\mathsf{S}_{k,k}$ walks are also expanding. A priori there is no reason why this claim should hold since a general d-sized γ-HDX may have much fewer hyperedges ($O_d(n)$ versus $\binom{n}{d}$ in the complete $\Delta_d(n)$). Fortunately, it turns out that this claim is indeed true (up to $O_k(\gamma)$ errors).

[2]The precise parameters of the complete complex $\Delta_d(n)$ where $\mathsf{S}^{\Delta}_{k,k}$ lives will not be important. We only require that $\mathsf{S}^{\Delta}_{k,k}$ is well defined in the sense that $d \ge 2k$ and $n > d$.

[3]The precise eigenvalues are also well known, but singular values are enough in our analysis.

To prove Theorem V.1[4] we employ the beautiful expanding poset (EPoset) machinery of Dikstein et al. [DDFH18]. Before we delve into the full technical analysis, it might be instructive to see how Theorem V.1 is obtained from understanding the quadratic form $\langle \mathsf{S}_{k,k} f, f\rangle$ where $f \in C^k$.

First we informally recall the decomposition $C^k = \sum_{i=0}^{k} C^k_i$ from the EPoset machinery where C^k_i can be thought of as the space of approximate eigenfunctions of *degree* i of C^k (the precise definitions are deferred to V-B). In this decomposition, C^k_0 is defined as the space of constant functions of C^k.

We prove the stronger result that the $\mathsf{S}_{k,k}$ operators of any γ-HDX has an an approximate spectrum that only depends on k provided γ is small enough. More precisely, we prove Lemma V.6.

Lemma V.6 (Swap Quadratic Form)**.** *Let $f=\sum_{i=0}^{k} f_i$ with $f_i \in C^k_i$. Suppose $X(\le d)$ is a γ-HDX with $d \ge 2k$. If $\gamma \le \varepsilon\left(64k^{k+4}2^{3k+1}\right)^{-1}$, then*

$$\langle \mathsf{S}_{k,k} f, f\rangle \;=\; \sum_{i=0}^{k} \lambda_k(i)\cdot\langle f_i, f_i\rangle \;\pm\; \varepsilon,$$

where $\lambda_k(i)$ depends only on k and i, i.e., $\lambda_k(i)$ is an approximate eigenvalue of $\mathsf{S}_{k,k}$ associated to space C^k_i.

Remark V.7. *From Lemma V.6, it might seem that we are done since there exist approximate eigenvalues $\lambda_k(i)$ that only depend on k and i. However, giving an explicit expression for these approximate eigenvalues is tricky. For this reason, we rely on the expansion of Kneser graphs as will be clear later.*

Towards showing Lemma V.6, we introduce the notion of *balanced* operators which in particular captures canonical and swap walks and we show that the quadratic form expression of Lemma V.6 is a particular case of a general result for $\langle \mathsf{B}f, f\rangle$ where B is a general *balanced* operator. A *balanced* operator in C^k is any operator that can be obtained as linear combination of *pure balanced* operators, the later being operators that are a formal product of an equal number of up and down operators.

Lemma V.8 (General Quadratic Form)**.** *Let $\varepsilon \in (0,1)$ and let $\mathcal{Y} \subseteq \{\mathsf{Y} \mid \mathsf{Y}\colon C^k \to C^k\}$ be a collection of formal operators that are product of an equal number of up and down walks (i.e., pure balanced operators) not exceeding ℓ walks. Let $\mathsf{B}=\sum_{\mathsf{Y}\in\mathcal{Y}} \alpha^{\mathsf{Y}}\mathsf{Y}$ where $\alpha^{\mathsf{Y}} \in \mathbb{R}$ and let $f=\sum_{i=0}^{k} f_i$ with $f_i \in C^k_i$. If $\gamma \le \varepsilon\left(16k^{k+2}\ell^2\sum_{\mathsf{Y}\in\mathcal{Y}}|\alpha^{\mathsf{Y}}|\right)^{-1}$, then*

$$\langle \mathsf{B}f, f\rangle \;=\; \sum_{i=0}^{k}\left(\sum_{\mathsf{Y}\in\mathcal{Y}} \alpha^{\mathsf{Y}}\lambda^{\mathsf{Y}}_k(i)\right)\cdot\langle f_i, f_i\rangle \;\pm\; \varepsilon,$$

[4]Due to space constraints, we will omit some details. The full proof can be found in the full version of our paper [AJT19]

where $\lambda_k^{\mathsf{Y}}(i)$ depends only on the operators appearing in the formal expression of Y, i and k, i.e., $\lambda_k^{\mathsf{Y}}(i)$ is the approximate eigenvalue of Y associated to C_i^k.

Remark V.9. *Note that our result generalizes the analysis of [DDFH18] for expanding posets of HDXs which considered the particular case $\mathsf{B} = \mathsf{D}_{k+1}\mathsf{U}_k$. Moreover, their error term analysis treated all the parameters not depending on the number of vertices n as constants. In this work we make the dependence on the parameters explicit since this dependence is important in understanding the limits of our k-CSPs approximation scheme on HDXs. The beautiful EPoset machinery [DDFH18] is instrumental in our analysis.*

Now, we are ready to prove Theorem V.1. For convenience we restate it below.

Theorem V.10 (Swap Walk Spectral Bound (restatement of Theorem V.1))**.** *Let $X(\leq d)$ be a γ-HDX with $d \geq 2k$. For every $\sigma \in (0,1)$, if $\gamma \leq \sigma \cdot \left(64k^{k+4}2^{3k+1}\right)^{-1}$, then the second largest singular value $\sigma_2(\mathsf{S}_{k,k})$ of the swap operator $\mathsf{S}_{k,k}$ is*

$$\sigma_2(\mathsf{S}_{k,k}) \;\leq\; \sigma.$$

B. Expanding Posets and Balanced Operators

We state the definitions used in our technical proofs starting with γ-EPoset from [DDFH18].

Definition V.11 (γ-EPoset adapted from [DDFH18])**.** *A complex $X(\leq d)$ with operators $\mathsf{U}_0,\ldots,\mathsf{U}_{d-1}$, $\mathsf{D}_1,\ldots,\mathsf{D}_d$ is said to be a γ-EPoset [5] provided*

$$\left\|\mathsf{M}_i^+ - \mathsf{U}_{i-1}\mathsf{D}_i\right\|_{\text{op}} \leq \gamma, \tag{1}$$

for every $i = 1,\ldots,d-1$, where

$$\mathsf{M}_i^+ \;\coloneqq\; \frac{i+1}{i}\left(\mathsf{D}_{i+1}\mathsf{U}_i - \frac{1}{i+1}\mathsf{I}\right),$$

i.e., M_i^+ is the non-lazy version of the random walk $\mathsf{N}_{i,i}^{(1)} = \mathsf{D}_{i+1}\mathsf{U}_i$.

Definition V.11 can be directly used as an operational definition of high-dimension expansion as done in [DDFH18]. To us it is important that γ-HDXs are also γ-EPosets as established in Lemma V.12. In fact, these two notions are known to be closely related.

Lemma V.12 (From [DDFH18])**.** *Let X be a d-sized simplicial complex.*

- *If X is a γ-HDX, then X is a γ-EPoset.*
- *If X is a γ-EPoset, then X is a $3d\gamma$-HDX.*

[5] We tailor their general EPoset definition to HDXs. In fact, what they call γ-HDX we call γ-EPoset. Moreover, what they call γ-HD expander we call γ-HDX.

Naturally the complete complex $\Delta_d(n)$ is a γ-EPoset since it is a γ-HDX. Moreover, in this particular case γ vanishes as n grows.

Lemma V.13 (From [DDFH18])**.** *The complete complex $\Delta_d(n)$ is a γ-EPoset with $\gamma = O_d(1/n)$.*

Harmonic Analysis on Simplicial Complexes: The space C^k defined in Section II-B2 can be decomposed into subspaces C_i^k of functions of *degree* i for $0 \leq i \leq k$ where

$$C_i^k \;\coloneqq\; \{\mathsf{U}^{k-i}h_i \mid h_i \in H_i\},$$

with $H_i \;\coloneqq\; \ker(\mathsf{D}_i)$, and

$$C_0^k \;\coloneqq\; \{f\colon X(k) \to \mathbb{R} \mid \; f \text{ is constant}\}.$$

More precisely, we have the following.

Lemma V.14 (From [DDFH18])**.**

$$C^k \;=\; \sum_{i=0}^{k} C_i^k.$$

For convenience set $\vec{\delta} \in \mathbb{R}^{d-1}$ such that $\delta_i = 1/(i+1)$ for $i \in [d-1]$. It will also be convenient to work with the following equivalent version of Eq. (1)

$$\left\|\mathsf{D}_{i+1}\mathsf{U}_i - (1-\delta_i)\mathsf{U}_{i-1}\mathsf{D}_i - \delta_i\mathsf{I}\right\|_{\text{op}} \;\leq\; \frac{i}{i+1}\gamma. \tag{2}$$

Towards our goal of understanding quadratic forms of swap operators we study the approximate spectrum of operators of the form $\mathsf{Y} = \mathsf{Y}_\ell \ldots \mathsf{Y}_1$ where each Y_i is either an up or down operator, namely, Y is a generalized random walk of ℓ steps. We regard the expression $\mathsf{Y}_\ell \ldots \mathsf{Y}_1$ defining Y as a formal product.

Definition V.15 (Pure Balanced Operator)**.** *We call $\mathsf{Y}\colon C^k \to C^k$ a pure balanced operator if Y can be defined as product $\mathsf{Y}_\ell \ldots \mathsf{Y}_1$ [6] where each Y_i is either an up or down operator. When we say that the spectrum of Y depends on Y we mean that it depends on k and on the formal expression $\mathsf{Y}_\ell \ldots \mathsf{Y}_1$ (i.e., pattern of up and down operators).*

Remark V.16. *By definition canonical walks $\mathsf{N}_{k,k}^{(u)}$ are pure balanced operators.*

Taking linear combinations of *pure balanced* operators leads to the notion of *balanced* operators.

Definition V.17 (Balanced Operator)**.** *We call $\mathsf{B}\colon C^k \to C^k$ a balanced operator provided there exists a set of pure balanced operators $\mathcal{Y}$ such that*

$$\mathsf{B} \;=\; \sum_{\mathsf{Y}\in\mathcal{Y}} \alpha^{\mathsf{Y}} \cdot \mathsf{Y},$$

where $\alpha^{\mathsf{Y}} \in \mathbb{R}$.

[6] For the analysis it is convenient to order the indices appearing in $\mathsf{Y}_\ell \ldots \mathsf{Y}_1$ in decreasing order from left to right.

Remark V.18. *Corollary IV.13 establishes that $\mathsf{S}_{k,k}^{(u)}$ are balanced operators. In particular, $\mathsf{S}_{k,k}$ is a balanced operator.*

It turns out that at a more crude level the behavior of Y is governed by how the number of up operators compares to the number of down operators. For this reason, it is convenient to define $\mathcal{U}(\mathsf{Y}) = \{\mathsf{Y}_i \mid \mathsf{Y}_i \text{ is an up operator}\}$ and $\mathcal{D}(\mathsf{Y}) = \{\mathsf{Y}_i \mid \mathsf{Y}_i \text{ is a down operator}\}$ where Y is a *pure balanced* operator. When Y is clear in the context we use $\mathcal{U} = \mathcal{U}(\mathsf{Y})$ and $\mathcal{D} = \mathcal{D}(\mathsf{Y})$.

Henceforth we assume $h_i \in H_i = \ker(\mathsf{D}_i)$, $f_i \in C_i^k$ and $g \in C^k$. This convention will make the statements of the technical results of Section V-C cleaner.

C. Quadratic Forms over Balanced Operators

Now we establish all the technical results leading to and including the analysis of quadratic forms over *balanced operators*. By considering this general class of operators our analysis generalizes the analysis given in [DDFH18]. At the same time we refine their error terms analysis by making the dependence on the EPoset parameters explicit. Recall that an explicit dependence on these parameters is important in understanding the limits of our k-CSP approximation scheme.

Lemma V.19 (General Quadratic Form (restatement of Lemma V.8))**.** *Let $\varepsilon \in (0,1)$ and let $\mathcal{Y} \subseteq \{\mathsf{Y} \mid \mathsf{Y}\colon C^k \to C^k\}$ be a collection of formal operators that are product of an equal number of up and down walks (i.e., pure balanced operators) not exceeding ℓ walks. Let $\mathsf{B} = \sum_{\mathsf{Y}\in\mathcal{Y}} \alpha^{\mathsf{Y}}\mathsf{Y}$ where $\alpha^{\mathsf{Y}} \in \mathbb{R}$ and let $f = \sum_{i=0}^{k} f_i$ with $f_i \in C_i^k$. If $\gamma \le \varepsilon\left(16k^{k+2}\ell^2\sum_{\mathsf{Y}\in\mathcal{Y}}|\alpha^{\mathsf{Y}}|\right)^{-1}$, then*

$$\langle \mathsf{B}f, f\rangle \;=\; \sum_{i=0}^{k}\left(\sum_{\mathsf{Y}\in\mathcal{Y}}\alpha^{\mathsf{Y}}\lambda_k^{\mathsf{Y}}(i)\right)\cdot\langle f_i, f_i\rangle \;\pm\; \varepsilon,$$

where $\lambda_k^{\mathsf{Y}}(i)$ depends only on the operators appearing in the formal expression of Y, i and k, i.e., $\lambda_k^{\mathsf{Y}}(i)$ is the approximate eigenvalue of Y associated to C_i^k.

Since swap walks are *balanced operators*, we will deduce the following.

Lemma V.20 (Swap Quadratic Form (restatement of Lemma V.6))**.** *Let $f = \sum_{i=0}^{k} f_i$ with $f_i \in C_i^k$. Suppose $X(\le d)$ is a γ-HDX with $d \ge 2k$. If $\gamma \le \varepsilon\left(64k^{k+4}2^{3k+1}\right)^{-1}$, then*

$$\langle \mathsf{S}_{k,k}f, f\rangle \;=\; \sum_{i=0}^{k}\lambda_k(i)\cdot\langle f_i, f_i\rangle \;\pm\; \varepsilon,$$

where $\lambda_k(i)$ depends on only on k an i, i.e., $\lambda_k(i)$ is an approximate eigenvalue of $\mathsf{S}_{k,k}$ associated to space C_i^k.

The next result, Lemma V.21, (implicit in [DDFH18]) will be key in establishing that the spectral structure of γ-EPosets is fully determined by the parameters in $\vec{\delta}$ provided γ is small enough. Note that the Eposet Definition V.11 provides a "calculus" for rearranging a single pair of up and down DU. The next result treats the more general case of $\mathsf{DU}\cdots\mathsf{U}$.

Lemma V.21 (Structure Lemma)**.** *Suppose $|\mathcal{D}| = 1$. Let $\mathsf{Y}_c \in \mathcal{D}$ be the unique down operator in $\mathsf{Y}_\ell \ldots \mathsf{Y}_1$. If $\|\mathsf{A}\|_{\text{op}} \le 1$, then*

$$\langle \mathsf{A}\mathsf{Y}_\ell \ldots \mathsf{Y}_1 h_i, g\rangle \;=\; \begin{cases} 0 & \text{if } \ell = 1 \text{ or } c = 1 \\ Q_{c,i}(\vec{\delta})\cdot\langle \mathsf{A}\mathsf{U}^{\ell-2}h_i, g\rangle \pm (c-1)\cdot\gamma\,\|h_i\|\,\|g\| & \text{otherwise,} \end{cases}$$

where $Q_{c,i}$ is a polynomial in the variables $\vec{\delta}$ depending on c, i such that $Q_{c,i}(\vec{\delta}) \le 1$.

With Lemma V.21 we are close to recover the approximate spectrum of $\mathsf{D}_{k+1}\mathsf{U}_k$ from [DDFH18]. However, in our application we will need to analyze more general operators, namely, *pure balanced* and *balanced* operators.

Lemma V.22 (Refinement of [DDFH18])**.** *If $\|\mathsf{A}\|_{\text{op}} \le 1$, then*

$$\langle \mathsf{A}\mathsf{D}_{k+1}\mathsf{U}_k f_i, g\rangle \;=\; \lambda_i\cdot\langle \mathsf{A}f_i, g\rangle \;\pm\; (k-i+1)\cdot\gamma\,\|h_i\|\,\|g\|\,,$$

where $\lambda_i = Q_{k-i+2,i}(\vec{\delta})$.

Then powers of the operator $\mathsf{D}_{k+1}\mathsf{U}_k$ behave as expected.

Lemma V.23 (Exponentiation Lemma)**.**

$$\langle (\mathsf{D}_{k+1}\mathsf{U}_k)^s f_i, f_i\rangle \;=\; \lambda_i^s\cdot\|f_i\|^2 \;\pm\; s\cdot(k-i+1)\cdot\gamma\,\|h_i\|\,\|f_i\|\,,$$

where λ_i is given in Lemma V.22.

In case $|\mathcal{D}| > |\mathcal{U}|$, $\mathsf{Y}\colon C^i \to C^j$ is an operator whose kernel approximately contains $\ker(\mathsf{D}_i)$ as the following lemma makes precise.

Lemma V.24 (Refinement of [DDFH18])**.** *If $|\mathcal{D}| > |\mathcal{U}|$ and $h_i \in \ker(\mathsf{D}_i)$, then*

$$\langle \mathsf{A}\mathsf{Y}_\ell \ldots \mathsf{Y}_1 h_i, g\rangle \;=\; \pm\,\ell^2\cdot\gamma\,\|h_i\|\,\|g\|\,,$$

provided $\|\mathsf{A}\|_{\text{op}} \le 1$.

We turn to an important particular case of $|\mathcal{D}| = |\mathcal{U}|$, namely, the canonical walks. We show that $\mathsf{N}_{k,k}^{(u)}$ is approximately a polynomial in the operator $\mathsf{D}_{k+1}\mathsf{U}_k$. As a warm up consider the case $\mathsf{N}_{k,k}^{(2)} = \mathsf{D}_{k+1}\mathsf{D}_{k+2}\mathsf{U}_{k+1}\mathsf{U}_k$. Using the Eq. (2), we get

$$\begin{aligned} \mathsf{N}_{k,k}^{(2)} \;&\approx\; (1-\delta_{k+1})\cdot\mathsf{D}_{k+1}\mathsf{U}_k\mathsf{D}_{k+1}\mathsf{U}_k \;+\; \delta_{k+1}\cdot\mathsf{D}_{k+1}\mathsf{U}_k \\ &=\; (1-\delta_{k+1})\cdot(\mathsf{D}_{k+1}\mathsf{U}_k)^2 \;+\; \delta_{k+1}\cdot\mathsf{D}_{k+1}\mathsf{U}_k. \end{aligned}$$

Inspecting this polynomial more carefully we see that that its coefficients form a probability distribution. This property holds in general as the following Lemma V.25 shows. This gives an alternative (approximate) random walk interpretation of $\mathsf{N}_{k,k}^{(u)}$ as the walk that first selects the power s

according to the distribution encoded in the polynomial and then moves according to $(\mathsf{D}_{k+1}\mathsf{U}_k)^s$.

Lemma V.25 (Canonical Polynomials)**.** *For $k, u \geq 0$ there exists a degree u univariate polynomial $F^N_{u,k,\vec{\delta}}$ depending only on $u, k, \vec{\delta}$ such that*

$$\left\| \mathsf{N}^{(u)}_{k,k} - F^N_{u,k,\vec{\delta}}(\mathsf{D}_{k+1}\mathsf{U}_k) \right\|_{\text{op}} \;\leq\; (u-1)^2 \cdot \gamma.$$

Moreover, the coefficients of this polynomial form a probability distribution, i.e., $F^N_{u,k,\vec{\delta}}(x) = \sum_{i=0}^{u} c_i x^i$ where $\sum_{i=0}^{u} c_i = 1$ and $c_i \geq 0$ for $i = 0, \ldots, u$.

Remark V.26. *Having a polynomial expression $F^N_{u,k,\vec{\delta}}(\mathsf{D}_{k+1}\mathsf{U}_k) \approx \mathsf{N}^{(u)}_{k,k}$ and knowing that $\mathsf{S}_{k,k}$ can be written as linear combination of canonical walks, we could deduce that $\mathsf{S}_{k,k}$ is also approximately a polynomial in $\mathsf{D}_{k+1}\mathsf{U}_k$. Using an error refined version of the Lemma V.23 (showing that exponentiation of $\mathsf{D}_{k+1}\mathsf{U}_k$ behaves naturally), we could deduce the approximate spectrum of $\mathsf{S}_{k,k}$. We avoid this approach since it analysis introduces unnecessary error terms and we can understand quadratic forms of pure balanced operators directly.*

Remark V.27. *The canonical polynomial $F^N_{u,k,\vec{\delta}}(\mathsf{D}_{k+1}\mathsf{U}_k)$ is used later in the error analysis that relates the norms $\|h_i\|$ and $\|f_i\|$ (Lemma V.30).*

Now we consider Y where $|\mathcal{D}| = |\mathcal{U}|$ in full generality. We show how the quadratic form of Y behaves in terms of the approximate eigenspace decomposition $C^k = \sum_{i=0}^{k} C^k_i$.

Lemma V.28 (Pure Balanced Walks)**.** *Suppose $\mathsf{Y} = \mathsf{Y}_\ell \ldots \mathsf{Y}_1$ is a product of an equal number of up and down operators, i.e., $|\mathcal{D}| = |\mathcal{U}|$. Then for $f_i \in C^k_i$*

$$\langle \mathsf{Y} f_i, f_i \rangle \;=\; \lambda^{\mathsf{Y}}_{k,i} \cdot \langle f_i, f_i \rangle \pm \gamma \cdot (\ell^2 + \ell(k-i-1)) \|h_i\| \, \|f_i\| \,,$$

where $\lambda^{\mathsf{Y}}_{k,i}$ is an approximate eigenvalue depending only on Y, k and i.

To understand all errors in the analysis in Lemma V.28we need to use the approximate orthogonality of f_i and f_j for $i \neq j$ from [DDFH18] in more detail. We start with the following bound in terms of h_i, h_j.

Lemma V.29 (Refinement of [DDFH18])**.** *For $i \neq j$,*

$$\langle f_i, f_j \rangle \;=\; \pm (2k-i-j)^2 \cdot \gamma \, \|h_i\| \, \|h_j\| \,.$$

To give a bound for Lemma V.29 only in terms of the eigenfunction norms $\|f_i\|$ and not in terms of $\|h_i\|$, we need to understand how the norms of h_i and f_i are related.

Lemma V.30 (Refinement of [DDFH18])**.** *Let $\eta_{k,i} = (k-i)^2 + 1$ and let $\beta_i = \sqrt{\left| F^N_{k-i,i,\vec{\delta}}(\delta_i) \pm \gamma \cdot \eta_{k,i} \right|}$ where $F^N_{k-i,k,\vec{\delta}}$ is a canonical polynomial of degree $k-i$ from Lemma V.25. Then*

$$\langle f_i, f_i \rangle \;=\; \beta_i^2 \cdot \langle h_i, h_i \rangle.$$

Let $\theta_{k,i} = (i+1)^{k-i}$. Furthermore, if $\gamma \leq 1/(2 \cdot \eta_{k,i} \cdot \theta_{k,i})$, then $\beta_i \geq \frac{1}{2\theta_{k,i}}$.

Now, we can state the approximate orthogonality Lemma V.31 in terms of the eigenfunction norms.

Lemma V.31 (Approximate Orthogonality (refinement of [DDFH18]))**.** *Let $\eta_{k,s}, \theta_{k,s}, \beta_s$ for $s \in \{i, j\}$ be given as in Lemma V.30. If $i \neq j$ and $\beta_i, \beta_j > 0$, then*

$$\langle f_i, f_j \rangle \;=\; \pm \; \frac{\gamma \cdot (2k-i-j)^2}{\beta_i \beta_j} \, \|f_i\| \, \|f_j\| \,.$$

Furthermore, if $\gamma \leq \min\left(1/(2 \cdot \eta_{k,i} \cdot \theta_{k,i}), 1/(2 \cdot \eta_{k,j} \cdot \theta_{k,j})\right)$, then $\beta_i, \beta_j > 0$ and

$$\langle f_i, f_j \rangle \;=\; \pm \, \gamma \cdot \theta_{k,i} \cdot \theta_{k,j} \cdot (2k-i-j)^2 \, \|f_i\| \, \|f_j\| \,.$$

We generalize the quadratic form of Lemma V.28 to linear combinations of general *pure balanced* operators Y, namely, to *balanced* operators.

Lemma V.32 (General Quadratic Form (restatement of Lemma V.8))**.** *Let $\varepsilon \in (0,1)$ and let $\mathcal{Y} \subseteq \{\mathsf{Y} \mid \mathsf{Y} \colon C^k \to C^k\}$ be a collection of formal operators that are product of an equal number of up and down walks (i.e., pure balanced operators) not exceeding ℓ walks. Let $\mathsf{B} = \sum_{\mathsf{Y} \in \mathcal{Y}} \alpha^{\mathsf{Y}} \mathsf{Y}$ where $\alpha^{\mathsf{Y}} \in \mathbb{R}$ and let $f = \sum_{i=0}^{k} f_i$ with $f_i \in C^k_i$. If $\gamma \leq \varepsilon \left(16 k^{k+2} \ell^2 \sum_{\mathsf{Y} \in \mathcal{Y}} |\alpha^{\mathsf{Y}}|\right)^{-1}$, then*

$$\langle \mathsf{B} f, f \rangle \;=\; \sum_{i=0}^{k} \left(\sum_{\mathsf{Y} \in \mathcal{Y}} \alpha^{\mathsf{Y}} \lambda^{\mathsf{Y}}_k(i) \right) \cdot \langle f_i, f_i \rangle \;\pm\; \varepsilon,$$

where $\lambda^{\mathsf{Y}}_k(i)$ depends only on the operators appearing in the formal expression of Y, i and k, i.e., $\lambda^{\mathsf{Y}}_k(i)$ is the approximate eigenvalue of Y associated to C^k_i.

We instantiate Lemma V.31 for swap walks with their specific parameters. First, we introduce some notation. Using Corollary IV.13, we have

$$\mathsf{S}_{k,k} \;=\; \sum_{j=0}^{k} (-1)^{k-j} \cdot \binom{k+j}{k} \cdot \binom{k}{j} \cdot \mathsf{N}^{(j)}_{k,k} \;=\; \sum_{j=0}^{k} \alpha_j \cdot \mathsf{N}^{(j)}_{k,k},$$

where $\alpha_j = (-1)^{k-j} \cdot \binom{k+j}{k} \cdot \binom{k}{j}$.

Finally, we have all the pieces to prove Lemma V.6 restated below.

Lemma V.33 (Swap Quadratic Form (restatement of Lemma V.6))**.** *Let $f = \sum_{i=0}^{k} f_i$ with $f_i \in C^k_i$. Suppose $X(\leq d)$ is a γ-HDX with $d \geq 2k$. If $\gamma \leq \varepsilon \left(64 k^{k+4} 2^{3k+1}\right)^{-1}$, then*

$$\langle \mathsf{S}_{k,k} f, f \rangle \;=\; \sum_{i=0}^{k} \lambda_k(i) \cdot \langle f_i, f_i \rangle \;\pm\; \varepsilon,$$

where $\lambda_k(i)$ depends on only on k an i, i.e., $\lambda_k(i)$ is an approximate eigenvalue of $\mathsf{S}_{k,k}$ associated to space C^k_i.

D. Rectangular Swap Walks $\mathsf{S}_{k,l}$

We turn to the spectral analysis of rectangular swap walks, i.e., $\mathsf{S}_{k,l}$ where $k \neq l$. Recall that to bound $\sigma_2(\mathsf{S}_{k,k})$ in Section V-A we proved that the spectrum of $\mathsf{S}_{k,k}$ for a γ-HDX is close to the spectrum of $\mathsf{S}^{\Delta}_{k,k}$ using the analysis of quadratic forms over *balanced* operators from Section V-C. Then we appealed to the fact that $\mathsf{S}^{\Delta}_{k,k}$ is expanding since it is the walk operator of the well known Kneser graph. In this rectangular case, we do not have a classical result establishing that $\mathsf{S}^{\Delta}_{k,l}$ is expanding, but we were able to establish it Lemma V.34.

Lemma V.34. *Let $d \geq k+l$ and $\Delta_d(n)$ be the complete complex. The second largest singular value $\sigma_2(\mathsf{S}^{\Delta}_{k,l})$ of the swap operator $\mathsf{S}^{\Delta}_{k,l}$ on $\Delta_d(n)$ is*

$$\sigma_2(\mathsf{S}^{\Delta}_{k,l}) \;\leq\; \max\left(\frac{k}{n-k}, \frac{l}{n-l}\right),$$

provided $n \geq M_{k,l}$ where $M_{k,l} \in \mathbb{N}$ only depends on k and l.

Towards proving Lemma V.34 we first introduce a generalization of Kneser graphs which we denote *bipartite Kneser* graphs defined as follows.

Definition V.35 (General Bipartite Kneser Graph). *Let $X(\leq d)$ where $d \geq k+l$. We denote by $K^X(n,k,l)$ the bipartite graph on (vertex) partition $(X(k), X(l))$ where $\mathfrak{s} \in X(k)$ is adjacent to $\mathfrak{t} \in X(l)$ if and only if $\mathfrak{s} \cap \mathfrak{t}$ is empty. We also refer to graphs of the form $K^X(n,k,l)$ as bipartite Kneser graphs.*

It will be convenient to distinguish *bipartite Kneser* graphs coming from general γ-HDX and the complete complex $\Delta_d(n)$.

Definition V.36 (Complete Bipartite Kneser Graph). *Let $X(\leq d)$ where $d \geq k+l$. If X is the complete complex, i.e., $X = \Delta_d(n)$, then we denote $K^X(n,k,l)$ as simply as $K(n,k,l)$ and we refer to it as complete bipartite Kneser.*

We obtain the spectra of *bipartite Kneser* graphs generalizing [7] the classical result of Fact V.4. More precisely, we prove Lemma V.37.

Lemma V.37 (Bipartite Kneser Spectrum). *The non-zero eigenvalues of the (normalized) walk operator of $K(n,k,l)$ are $\pm\lambda_i$ where*

$$\lambda_i \;=\; \frac{\binom{n-k-i}{l-i}\binom{n-l-i}{k-i}}{\binom{n-k}{l}\binom{n-l}{k}},$$

for $i = 0, \ldots, \min(k,l)$.

[7]Note that the singular values of $K(n,k)$ can be deduced from the bipartite case.

Now the proof follows a similar strategy to the $\mathsf{S}_{k,k}$, namely, we analyze quadratic forms over $\mathsf{S}_{k,k}$ using the results from Section V-C

Let $X(\leq d)$ where $d \geq k+l$. Let $\mathsf{A}_{k,l}$ be the (normalized) walk operator of $K^X(n,k,l)$, i.e.,

$$\mathsf{A}_{k,l} \;=\; \begin{pmatrix} 0 & \mathsf{S}^{(l)}_{k,l} \\ \left(\mathsf{S}^{(l)}_{k,l}\right)^{\dagger} & 0 \end{pmatrix}.$$

To determine the spectrum of $\mathsf{A}_{k,l}$ it is enough to consider the spectrum of $\mathsf{B} = \mathsf{S}^{(l)}_{k,l}\left(\mathsf{S}^{(l)}_{k,l}\right)^{\dagger}$. Using Corollary IV.13, we have

$$\begin{aligned} \mathsf{B} \;&=\; \left(\sum_{j=0}^{l}(-1)^{l-j}\binom{k+j}{l}\cdot\binom{l}{j}\cdot\mathsf{N}^{(j)}_{k,l}\right) \\ &\qquad \left(\sum_{j'=0}^{l}(-1)^{l-j'}\binom{k+j'}{l}\cdot\binom{l}{j'}\cdot\left(\mathsf{N}^{(j')}_{k,l}\right)^{\dagger}\right) \\ &=\; \sum_{j,j'=0}^{l}\alpha_{k,l,j,j'}\mathsf{N}^{(j)}_{k,l}\mathsf{N}^{(j'+k-l)}_{l,k}, \end{aligned}$$

for some coefficients $\alpha_{k,l,j,j'}$ depending only on k, l, i, j and j'. Since we have not yet used any specific property of HDXs, these coefficients are the same for the complete complex and general HDXs.

Lemma V.38. *Let $X(\leq d)$ be a γ-HDX with $d \geq k+l$. Let $f = \sum_{i=0}^{k} f_i$ with $f_i \in C^k_i$. For $\varepsilon \in (0,1)$, if $\gamma \leq \varepsilon\left(64k^{k+2}\ell^2 2^{2k+4l+2}\right)^{-1}$, then*

$$\langle \mathsf{B}f, f\rangle \;=\; \sum_{i=0}^{k}\left(\sum_{j,j'=0}^{l}\alpha_{k,l,j,j'}\lambda_{k,l,j,j'}(i)\right)\cdot\langle f_i, f_i\rangle \;+\; \varepsilon,$$

where $\lambda_{k,l,j,j'}(i)$ is the approximate eigenvalues of $\mathsf{N}^{(j)}_{k,l}\mathsf{N}^{(j'+k-l)}_{l,k}$ corresponding to space C^k_i. Furthermore, $\lambda_{k,l,j,j'}(i)$ depends only on k, l, i, j and j'.

Let B and B^{Δ} stand for the B operator for general γ-HDX and the complete complex, respectively.

Lemma V.39. *Suppose $X(\leq d)$ is a γ-HDX with $d \geq k+l$. For $\varepsilon \in (0,1)$, if $\gamma \leq \varepsilon^2\left(64k^{k+2}\ell^2 2^{2k+4l+2}\right)^{-1}$, then the second largest singular value $\sigma_2(\mathsf{B})$ of B is*

$$\sigma_2(\mathsf{B}) \;\leq\; \varepsilon^2.$$

Furthermore, the second largest non-trivial eigenvalue $\lambda(\mathsf{A}_{k,l})$ of the walk matrix of $K(n,k,l)$ is

$$\lambda(\mathsf{A}_{k,l}) \;\leq\; \varepsilon.$$

Now the proof of Theorem V.2 follows. For convenience, we restate it.

Theorem V.40 (Rectangular Swap Walk Spectral Bound (restatement of Theorem V.2)). *Suppose $X(\leq d)$ is a γ-HDX with $d \geq k+l$ and $k \leq l$. For $\sigma \in (0,1)$, if*

$\gamma \le \sigma^2 \cdot \left(64k^{k+2}\ell^2 2^{2k+4l+2}\right)^{-1}$, *then the largest non-trivial singular value* $\sigma_2(\mathsf{S}_{k,l})$ *of the swap operator* $\mathsf{S}_{k,l}$ *is*

$$\sigma_2(\mathsf{S}_{k,l}) \;\le\; \sigma.$$

E. Bipartite Kneser Graphs - Complete Complex

Now we determine the spectrum of the *complete bipartite Kneser* graph $K(n,k,l)$. More precisely, we prove the following.

Lemma V.41 (Bipartite Kneser Spectrum (restatement of Lemma V.37))**.** *The non-zero eigenvalues of the normalized walk operator of* $K(n,k,l)$ *are* $\pm\lambda_i$ *where*

$$\lambda_i \;=\; \frac{\binom{n-k-i}{l-i}\binom{n-l-i}{k-i}}{\binom{n-k}{l}\binom{n-l}{k}},$$

for $i = 0, \dots, \min(k,l)$.

Henceforth, set $X = \Delta_d(n)$. To prove Lemma V.37 we work with the natural rectangular matrix associated with $K(n,k,l)$, namely, the matrix $W \in \mathbb{R}^{X(k)\times X(l)}$ such that

$$\mathsf{W}(\mathfrak{s},\mathfrak{t}) \;=\; \mathbb{1}_{[\mathfrak{s}\cap\mathfrak{t}=\emptyset]}$$

for every $\mathfrak{s} \in X(k)$ and $\mathfrak{t} \in X(l)$.

Observe that the entries of $\mathsf{W}\mathsf{W}^\top$ and $\mathsf{W}^\top\mathsf{W}$ only depend on the size of the intersection of the sets indexing the row and columns. Hence, these matrices belong to the Johnson scheme [GM15] $J(n,k)$ and $J(n,l)$, respectively. Moreover, the left and right singular vectors of W are eigenvectors of these schemes.

We adopt the eigenvectors used in Filmus' work [Fil16], i.e., natural basis vectors coming from some irreducible representation of S_n (see [Sag13]). First we introduce some notation. Let $\mu = (n-i,i)$ be a partition of n and let τ_μ be a standard tableau of shape μ. Suppose the first row τ_μ contains $a_1 < \dots < a_{n-i}$ whereas the second contains $b_1 < \dots < b_i$. To τ_μ we associate the function $\varphi_{\tau_\mu} \in \mathbb{R}^{\binom{[n]}{k}}$ as follows

$$\varphi_{\tau_\mu} \;=\; (\mathbb{1}_{a_1} - \mathbb{1}_{b_1})\dots(\mathbb{1}_{a_i} - \mathbb{1}_{b_i}),$$

where $\mathbb{1}_a \in \mathbb{R}^{\binom{n}{k}}$ is the containment indicator of element a, i.e., $\mathbb{1}_a(\mathfrak{s}) = 1$ if and only if $a \in \mathfrak{s}$. Filmus proved that

$$\left\{\varphi_{\tau_\mu} \mid 0 \le i \le k, \mu \vdash (n-i,i), \tau_\mu \text{ standard}\right\}$$

is an eigenbasis of $\mathcal{J}(n,k)$. We abuse the notation by considering φ_{τ_μ} as both a function in $\mathbb{R}^{\binom{n}{k}}$ and $\mathbb{R}^{\binom{n}{l}}$ as long as these functions are well defined.

Claim V.42. *If* $\mu = (n-i,i)$ *and* $k,l \ge i$*, then*

$$\mathsf{W}\varphi_{\tau_\mu} \;=\; (-1)^i \cdot \binom{n-k-i}{l-i} \cdot \varphi_{\tau_\mu}.$$

Since we are working with singular vectors, we need to be careful with their normalization when deriving the singular values. We stress that the norm of φ_{τ_μ} depends on the space where φ_{τ_μ} lies.

Claim V.43. *If* $\mu = (n-i,i)$ *and* $\varphi_{\tau_\mu} \in \mathbb{R}^{\binom{n}{k}}$*, then*

$$\left\|\varphi_{\tau_\mu}\right\|_2 \;=\; \sqrt{2^i\binom{n-2i}{k-i}}.$$

Now the singular values of W follow.

Corollary V.44 (Singular Values)**.** *The singular values of* W *are*

$$\sigma_i \;=\; \binom{n-k-i}{l-i} \cdot \frac{\left\|\varphi^k_{\tau_\mu}\right\|_2}{\left\|\varphi^l_{\tau_\mu}\right\|_2},$$

for $i = 0, \dots, \min(k,l)$.

Note that for $k = l$ we recover the well know result of Fact V.4.

Finally we compute the eigenvalues of the bipartite graph $K(n,k,l)$. Let $\mathsf{A}_{n,k,l}$ be its normalized adjacency matrix, i.e.,

$$\mathsf{A}_{n,k,l} \;=\; \begin{pmatrix} 0 & \frac{1}{\binom{n-k}{l}}\mathsf{W} \\ \frac{1}{\binom{n-l}{k}}\mathsf{W}^\top & 0 \end{pmatrix}.$$

Lemma V.45 (Bipartite Kneser Spectrum (restatement of Lemma V.37))**.** *The non-zero eigenvalues of the normalized walk operator of* $K(n,k,l)$ *are* $\pm\lambda_i$ *where*

$$\lambda_i \;=\; \frac{\binom{n-k-i}{l-i}\binom{n-l-i}{k-i}}{\binom{n-k}{l}\binom{n-l}{k}},$$

for $i = 0, \dots, \min(k,l)$.

VI. Approximating Max-k-CSP

In the following, we will show that k-CSP instances $\mathfrak{I}$ whose constraint complex $X_{\mathfrak{I}}(\le k)$ is a suitable expander admit an efficient approximation algorithm. We will assume throughout that $X_{\mathfrak{I}}(1) = [n]$, and drop the subscript $\mathfrak{I}$.

This was shown for 2-CSPs in [BRS11]. In extending this result to k-CSPs we will rely on a central Lemma of their paper. Before, we explain our algorithm we give a basic outline of our idea:

We will work with the SDP relaxation for the k-CSP problem given by L-levels of SoS hierarchy, as defined in Section II-D (for L to be specified later). This will give us an L-local PSD ensemble $\{\mathbf{Y}_1, \dots, \mathbf{Y}_n\}$, which attains some value $\mathsf{SDP}(\mathfrak{I}) \ge \mathsf{OPT}(\mathfrak{I})$. Since $\{\mathbf{Y}_1, \dots, \mathbf{Y}_n\}$, is a local PSD ensemble, and not necessarily a probability distribution, we cannot sample from it directly. Nevertheless, since $\{\mathbf{Y}_j\}$ will be actual probability distributions for all $j \in [n]$, one can independently sample $\eta_j \sim \{\mathbf{Y}_j\}$ and use $\eta = (\eta_1, \dots, \eta_n)$ as the assignment for the k-CSP instance $\mathfrak{I}$.

Unfortunately, while we know that the local distributions $\{\mathbf{Y}_{\mathfrak{a}}\}_{\mathfrak{a}\in X(k)}$ induced by $\{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$ will satisfy the constraints of $\mathfrak{I}$ with good probability, i.e.,

$$\mathop{\mathbb{E}}_{\mathfrak{a}\sim\Pi_k}\mathop{\mathbb{E}}_{\{\mathbf{Y}_{\mathfrak{a}}\}}\left[\mathbf{1}\underbrace{[\mathbf{Y}_a \text{ satisfies the constraint on } \mathfrak{a}]}_{\iff \mathbf{Y}_{\mathfrak{a}}\in\,\mathcal{C}_{\mathfrak{a}}}\right] = \mathsf{SDP}(\mathfrak{I}) \geq \mathsf{OPT}(\mathfrak{I}),$$

this might not be the case for the assignment η sampled as before. It might be that the random variables $\mathbf{Y}_{a_1},\ldots,\mathbf{Y}_{a_k}$ are highly correlated for $\mathfrak{a} \in X(k)$, i.e., $\mathbb{E}_{\mathfrak{a}\sim\Pi_k}\|\{\mathbf{Y}_{\mathfrak{a}}\} - \{\mathbf{Y}_{a_1}\}\cdots\{\mathbf{Y}_{a_k}\}\|_1$ is large. One strategy employed by [BRS11] to ensure that the quantity above is small, is making the local PSD ensemble $\{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$ be consistent with a randomly sampled partial assignment for a small subset of variables (q.v. Section II-D). We will show that this strategy is succesful if $X(\leq k)$ is a γ-HDX (for γ sufficiently small). Our final algorithm is Algorithm VI.1.

Algorithm VI.1 (Propagation Rounding Algorithm)**.**

Intput *An L-local PSD ensemble $\{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$, and a distribution Π on $X(\leq \ell)$.*

Output *A random assignment $\eta : [n] \to [q]$.*

1) *Choose $m \in \{1,\ldots,L/\ell\}$ uniformly at random.*
2) *Independently sample m ℓ-faces, $\mathfrak{s}_j \sim \Pi$ for $j = 1,\ldots,m$.*
3) *Write $S = \bigcup_{j=1}^m \mathfrak{s}_j$, for the set of the seed vertices.*
4) *Sample assignment $\eta_S : S \to [q]$ according to the local distribution, $\{\mathbf{Y}_S\}$.*
5) *Set $\mathbf{Y}' = \{\mathbf{Y}_1,\ldots\mathbf{Y}_n | \mathbf{Y}_S = \eta_S\}$, i.e. the local ensemble $\mathbf{Y}$ conditioned on agreeing with η_S.*
6) *For all $j \in [n]$, sample independently $\eta_j \sim \{\mathbf{Y}'_j\}$.*
7) *Output $\eta = (\eta_1,\ldots,\eta_n)$.*

In our setting, we will apply Algorithm VI.1 with the distribution Π_k and the L-local PSD ensemble $\{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$. Notice that in expectation, the marginals of $\mathbf{Y}'$ on faces $\mathfrak{a} \in X(k)$ – which are actual distributions – will agree with the marginals of $\mathbf{Y}$, i.e. $\mathbb{E}_{S,\eta_S}\mathbb{E}\,\mathbf{Y}'_{\mathfrak{a}} = \mathbb{E}\,\mathbf{Y}_{\mathfrak{a}}$. In particular, the approximation quality of Algorithm VI.1 will depend on the average correlation of $\mathbf{Y}'_{a_1},\ldots,\mathbf{Y}'_{a_k}$ on the constraints $\mathfrak{a} \in X(k)$, where $\mathbf{Y}'$ is the local PSD ensemble obtained at the end of the first phase of Algorithm VI.1.

In the case where $k = 2$, the following is known

Theorem VI.2 (Theorem 5.6 from [BRS11])**.** *Suppose a weighted undirected graph $G = ([n], E, \Pi_2)$ and an L-local PSD ensemble $Y = \{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$ are given. There exists absolute constants $c \geq 0$ and $C \geq 0$ satisfying the following: If $L \geq c\cdot\frac{q}{\varepsilon^4}$, $\mathrm{Supp}(\mathbf{Y}_i) \leq q$ for all $i \in V$, and $\lambda_2(G) \leq C\cdot\varepsilon^2/q^2$ then we have*

$$\mathop{\mathbb{E}}_{\{i,j\}\sim\Pi_2}\left\|\{\mathbf{Y}'_i,\mathbf{Y}'_j\} - \{\mathbf{Y}'_i\}\{\mathbf{Y}'_j\}\right\|_1 \leq \varepsilon,$$

where $\mathbf{Y}'$ is as defined in Algorithm VI.1 on the input of $\{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$ and Π_1.

To approximate k-CSPs well, we will show the following generalization of Theorem VI.2 for k-CSP instances $\mathfrak{I}$, whose constraint complex $X(\leq k)$ is γ-HDX, for γ sufficiently small.

Theorem VI.3. *Suppose a simplicial complex $X(\leq k)$ with $X(1) = [n]$ and an L-local PSD ensemble $\mathbf{Y} = \{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$ are given.*

There exists some universal constants $c' \geq 0$ and $C' \geq 0$ satisfying the following: If $L \geq c'\cdot(q^k\cdot k^5/\varepsilon^4)$, $\mathrm{Supp}(\mathbf{Y}_j) \leq q$ for all $j \in [n]$, and X is a γ-HDX for $\gamma \leq C'\cdot\varepsilon^4/(k^{8+k}\cdot 2^{6k}\cdot q^{2k})$. Then, we have

$$\mathop{\mathbb{E}}_{\mathfrak{a}\sim\Pi_k}\left\|\{\mathbf{Y}'_{\mathfrak{a}}\} - \{\mathbf{Y}'_{a_1}\}\cdots\{\mathbf{Y}'_{a_k}\}\right\|_1 \leq \varepsilon, \tag{3}$$

where $\mathbf{Y}'$ is as defined in Algorithm VI.1 on the input of $\{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$ and Π_k.

Indeed, using Theorem VI.3, it is straightforward to prove the following,

Corollary VI.4. *Suppose $\mathfrak{I}$ is a q-ary k-CSP instance whose constraint complex $X(\leq k)$ is a γ-HDX.*

There exists absolute constants $C' \geq 0$ and $c' \geq 0$, satisfying the following: If $\gamma \leq C'\cdot\varepsilon^4/(k^{8+k}\cdot 2^{6k}\cdot q^{2k})$, there is an algorithm that runs in time $n^{O(k^5\cdot q^{2k}\cdot\varepsilon^{-4})}$ based on $(\frac{c'\cdot k^5\cdot q^k}{\varepsilon^4})$-levels of SoS-hierarchy and Algorithm VI.1 that outputs a random assignment $\eta : [n] \to [q]$ that in expectation ensures $\mathsf{SAT}_{\mathfrak{I}}(\eta) = \mathsf{OPT}(\mathfrak{I}) - \varepsilon$.

Our proof of Theorem VI.3[8] will hinge on the fact that we can upper-bound the expected correlation of a face of large cardinality ℓ, in terms of expected correlation over faces of smaller cardinality and expected correlations along the edges of a swap graph. The swap graph here is defined as a weighted graph $G_{\ell_1,\ell_2} = (X(\ell_1) \sqcup X(\ell_2), E(\ell_1,\ell_2), w_{\ell_1,\ell_2})$, where

$$E(\ell_1,\ell_2) = \left\{\{\mathfrak{a},\mathfrak{b}\} : \begin{array}{l}\mathfrak{a}\in X(\ell_1), \mathfrak{b}\in X(\ell_2),\\ \text{and } \mathfrak{a}\sqcup\mathfrak{b}\in X(\ell_1+\ell_2)\end{array}\right\}.$$

We will assume $\ell_1 \geq \ell_2$, and if $\ell_1 = \ell_2$ we are going to identify the two copies of every vertex. We will endow $E(\ell_1,\ell_2)$ with the weight function,

$$w_{\ell_1,\ell_2}(\mathfrak{a},\mathfrak{b}) = \frac{\Pi_{\ell_1+\ell_2}(\mathfrak{a}\sqcup\mathfrak{b})}{\binom{\ell_1+\ell_2}{\ell_1}},$$

which can easily be verified to be a probability distribution on $E(\ell_1,\ell_2)$ Notice that in the case where $\ell_1 \neq \ell_2$ the random walk matrix of G_{ℓ_1,ℓ_2} is given by

$$\mathsf{A}_{\ell_1,\ell_2} = \begin{pmatrix} 0 & \mathsf{S}_{\ell_1,\ell_2}\\ \mathsf{S}^{\dagger}_{\ell_1,\ell_2} & 0\end{pmatrix},$$

[8] Due to space constraints, we will omit some details. The full proof can be found in the full version of our paper [AJT19].

and if $\ell_1 = \ell_2$ we have $\mathsf{A}_{\ell_1,\ell_1} = \mathsf{S}_{\ell_1,\ell_1}$. The stationary distribution of $\mathsf{A}_{\ell_1,\ell_2}$ is Π_{ℓ_1,ℓ_2} defined by,

$$\Pi_{\ell_1,\ell_2}(\mathfrak{b}) = \frac{\mathbf{1}[\mathfrak{b} \in X(\ell_1)]}{2} \cdot \Pi_{\ell_1}(\mathfrak{b}) + \frac{\mathbf{1}[\mathfrak{b} \in X(\ell_2)]}{2} \cdot \Pi_{\ell_2}(\mathfrak{b}). \quad (4)$$

When we write an expectation of $\boldsymbol{f}(\bullet,\bullet)$ over the edges in $E(\ell_1,\ell_2)$ with respect to w_{ℓ_1,ℓ_2}, it is important to note,

$$\begin{aligned} \mathop{\mathbb{E}}_{\{\mathfrak{s},\mathfrak{t}\}\sim w_{\ell_1,\ell_2}} [\boldsymbol{f}(\mathfrak{s},\mathfrak{t})] &= \sum_{\{\mathfrak{s},\mathfrak{t}\}\in E(\ell_1,\ell_2)} \frac{\Pi_{\ell_1+\ell_2}(\mathfrak{s}\sqcup\mathfrak{t}),}{\binom{\ell_1+\ell_2}{\ell_1}} \cdot \boldsymbol{f}(\mathfrak{s},\mathfrak{t}), \\ &= \frac{1}{\binom{\ell}{\ell_1}} \mathop{\mathbb{E}}_{\mathfrak{a}\sim\Pi_k} \left[\sum_{\{\mathfrak{s},\mathfrak{t}\}\sim\mathfrak{a}} \boldsymbol{f}(\mathfrak{s},\mathfrak{t}) \right], \quad (5) \end{aligned}$$

where sum within the expectation in the RHS runs over the $\binom{\ell_1+\ell_2}{\ell_1}$ possible ways of splitting $\mathfrak{a}$ into $\mathfrak{s} \sqcup \mathfrak{t}$ such that $\mathfrak{s} \in X(\ell_1)$ and $\mathfrak{t} \in X(\ell_2)$. When we are speaking about the spectral expansion of G_{ℓ_1,ℓ_2}, we will be speaking with regards to $\lambda_2(G_{\ell_1,\ell_2})$ and not with regards to $\sigma_2(G_{\ell_1,\ell_2})$.

Remark VI.5. *By simple linear algebra, we have*

$$\lambda_2(G_{\ell_1,\ell_2}) := \lambda_2(\mathsf{A}_{\ell_1,\ell_2}) \le \sigma_2(\mathsf{S}_{\ell_1,\ell_2}),$$

where we employ the notation $\lambda_2(\mathsf{M})$ to denote the second largest eigenvalue (signed) of the matrix M.

With this, we will show

Lemma VI.6 (Glorified Triangle Inequality)**.** *For a simplicial complex $X(\le k)$, $\ell_1 \ge \ell_2 \ge 0$, $\ell = \ell_1 + \ell_2$, $\ell \le k$, and an ℓ-local ensemble $\{\mathbf{Y}_1, \dots, \mathbf{Y}_n\}$, one has*

$$\begin{aligned} \mathop{\mathbb{E}}_{\mathfrak{a}\in\Pi_\ell} \left\| \{\mathbf{Y}_\mathfrak{a}\} - \prod_{i=1}^{\ell} \{\mathbf{Y}_{a_i}\} \right\|_1 &\le \mathop{\mathbb{E}}_{\{\mathfrak{s},\mathfrak{t}\}\sim w_{\ell_1,\ell_2}} \|\{\mathbf{Y}_\mathfrak{s}, \mathbf{Y}_\mathfrak{t}\} - \{\mathbf{Y}_\mathfrak{s}\}\{\mathbf{Y}_\mathfrak{t}\}\|_1 \\ &+ \mathop{\mathbb{E}}_{\mathfrak{s}\sim\Pi_{\ell_1}} \left\| \{\mathbf{Y}_\mathfrak{s}\} - \prod_{i=1}^{\ell_1} \{\mathbf{Y}_{s_i}\} \right\|_1 + \mathop{\mathbb{E}}_{\mathfrak{t}\sim\Pi_{\ell_2}} \left\| \{\mathbf{Y}_\mathfrak{t}\} - \prod_{i=1}^{\ell_2} \{\mathbf{Y}_{t_i}\} \right\|_1 \quad (6) \end{aligned}$$

One useful observation, is that by using Lemma VI.6 repeatedly, we can reduce the problem of bounding $\mathbb{E}_{\mathfrak{a}\in\Pi_\ell} \left\| \{\mathbf{Y}_\mathfrak{a}\} - \prod_{i=1}^{\ell} \{\mathbf{Y}_{a_i}\} \right\|_1$ to a problem of bounding

$$\mathop{\mathbb{E}}_{\{\mathfrak{s},\mathfrak{t}\}\sim w_{\ell_1,\ell_2}} \|\{\mathbf{Y}_\mathfrak{s}, \mathbf{Y}_\mathfrak{t}\} - \{\mathbf{Y}_\mathfrak{s}\}\{\mathbf{Y}_\mathfrak{t}\}\|_1,$$

for $\ell_1 + \ell_2 \le k$. Though it is not a direct implication, it is heavily suggested by Fact II.7 and Theorem VI.2, that if G_{ℓ_1,ℓ_2} is a good spectral expander, after an application of Algorithm VI.1 with our chosen parameters, we should be able to bound these expressions. Using a key lemma used from [BRS11], we will prove that this is indeed the case. The only thing we need to make sure after this point, is that the second eigenvalue $\lambda_2(G_{\ell_1,\ell_2})$ of the swap graphs G_{ℓ_1,ℓ_2} we will be using are small enough for our purposes. Indeed, our choice of γ in Theorem VI.3 and Corollary VI.4 is to make sure that the bound we get on $\lambda_2(G_{\ell_1,\ell_2})$ from Theorem V.2 (together with Remark VI.5) is good enough for our purposes.

A. Breaking Correlations for Expanding CSPs

Throughout this section, we will use the somewhat non-standard definition of variance introduced in [BRS11],

$$\mathrm{Var}[\mathbf{Y}_\mathfrak{a}] = \sum_{\eta\in[q]^\mathfrak{a}} \mathrm{Var}[\mathbf{1}[\mathbf{Y}_\mathfrak{a} = \eta]].$$

We will use the following central lemma from [BRS11] in our proof of Theorem VI.3:

Lemma VI.7 (Lemma 5.4 from [BRS11])**.** *Let $G = (V, E, \Pi_2)$ be a weighted graph, $\{\mathbf{Y}_1, \dots, \mathbf{Y}_n\}$ a local PSD ensemble, where we have $\mathrm{Supp}(\mathbf{Y}_i) \le q$ for every $i \in V$, and $q \ge 0$. Suppose $\varepsilon \ge 0$ is a lower bound on the expected statistical difference between independent and correlated sampling along the edges,i.e.,*

$$\varepsilon \le \mathop{\mathbb{E}}_{\{i,j\}\sim\Pi_2} \|\{\mathbf{Y}_{ij}\} - \{\mathbf{Y}_i\}\{\mathbf{Y}_j\}\|_1.$$

There exists absolute constants $c_0 \ge 0$ and $c_1 \ge 0$ that satisfy the following: If $\lambda_2(G) \le c_0 \cdot \frac{\varepsilon^2}{q^2}$. Then, conditioning on a random vertex decreases the variances,

$$\mathop{\mathbb{E}}_{i\sim\Pi_1} \mathop{\mathbb{E}}_{j\sim\Pi_2} \mathop{\mathbb{E}}_{\{\mathbf{Y}_j\}} \mathrm{Var}[\mathbf{Y}_i \mid \mathbf{Y}_j] \le \mathop{\mathbb{E}}_{i\sim\Pi_1} [\mathrm{Var}[\mathbf{Y}_i]] - c_1 \cdot \frac{\varepsilon^2}{q^2}.$$

For our applications, we will be instantiating Lemma VI.7 with G_{ℓ_1,ℓ_2} as G; and with the local PSD ensemble $\{\mathbf{Y}_\mathfrak{a}\}_{\mathfrak{a}\in X}$ that is obtained from $\{\mathbf{Y}_1, \dots, \mathbf{Y}_n\}$ (q.v. Fact II.7). For convenience, we will write the concrete instance of the Lemma that we will use,

Corollary VI.8. *Let $\ell_1 \ge \ell_2 \ge 0$ satisfying $\ell_1 + \ell_2 \le k$ be given parameters, and let G_{ℓ_1,ℓ_2} be the swap graph defined for a γ-HDX $X(\le k)$. Let $\{\mathbf{Y}_\mathfrak{a}\}_{\mathfrak{a}\in X}$ be a local PSD ensemble; satisfying $\mathrm{Supp}(Y_\mathfrak{a}) \le q^k$ for every $\mathfrak{a} \in X(\ell_1) \cup X(\ell_2)$ for some $q \ge 0$. Suppose $\varepsilon \ge 0$ satisfies,*

$$\frac{\varepsilon}{4k} \le \mathop{\mathbb{E}}_{\{\mathfrak{s},\mathfrak{t}\}\in w_{\ell_1,\ell_2}} \|\{\mathbf{Y}_{\mathfrak{s}\sqcup\mathfrak{t}}\} - \{\mathbf{Y}_\mathfrak{s}\}\{\mathbf{Y}_\mathfrak{t}\}\|_1.$$

There exists absolute constants $c_0 \ge 0$ and $c_2 \ge 0$ that satisfy the following: If $\lambda_2(G) \le c_0 \cdot (\varepsilon/(4k \cdot q^k))^2$. Then, conditioning on a random face $\mathfrak{a} \sim \Pi_{\ell_1,\ell_2}$ decreases the variances, i.e.

$$\begin{aligned} &\mathop{\mathbb{E}}_{\mathfrak{a},\mathfrak{b}\sim\Pi^2_{\ell_1,\ell_2}} \mathop{\mathbb{E}}_{\{\mathbf{Y}_\mathfrak{a}\}} [\mathrm{Var}[\mathbf{Y}_\mathfrak{b} \mid \mathbf{Y}_\mathfrak{a}]] \\ &\le \mathop{\mathbb{E}}_{\mathfrak{b}\sim\Pi_{\ell_1,\ell_2}} [\mathrm{Var}[\mathbf{Y}_\mathfrak{b}]] - c_2 \cdot \frac{\varepsilon^2}{16 \cdot k^2 \cdot q^{2k}}. \end{aligned}$$

Here, it can be verified that the expansion criterion presupposed by Lemma VI.7 is satisfied by Corollary VI.8 by Theorem V.2.

VII. High-Dimensional Threshold Rank

In [BRS11], Theorem VI.2 was proven for a more general class of graphs than expander graphs – namely, the class of low threshold rank graphs.

Definition VII.1 (Threshold Rank of Graphs (from [BRS11]))**.** *Let $G = (V, E, w)$ be a weighted graph on n vertices and A be its normalized random walk matrix. Suppose the eigenvalues of A are $1 = \lambda_1 \geq \cdots \geq \lambda_n$. Given a parameter $\tau \in (0, 1)$, we denote the threshold rank of G by $\operatorname{rank}_{\geq\tau}(\mathsf{A})$ (or $\operatorname{rank}_{\geq\tau}(G)$) and define it as*

$$\operatorname{rank}_{\geq\tau}(\mathsf{A}) := |\{i | \lambda_i \geq \tau\}| .$$

There [BRS11], the authors asked for the correct notion of threshold rank for k-CSPs. In this section, we give a candidate definition of low threshold rank motivated by our techniques.

To break k-wise correlations it is sufficient to assume that the involved swap graphs in the foregoing discussion are low threshold rank since this is enough to apply a version of Lemma VI.7, already described in the work of [BRS11].

Moreover, we have some flexibility as to which swap graphs to consider as long as they satisfy some splitting conditions. To define a swap graph it is enough to have a distributions on the hyperedges of a (constraint) hypergraph. Hence, the notion of swap graph is independent of high-dimensional expansion. HDXs are just an interesting family of objects for which the swap graphs are good expanders.

To capture the many ways of splitting the statistical distance over hyperedges into the statistical distance over the edges of swap graphs, we first define the following notion. We say that a binary tree $\mathcal{T}$ is a k-splitting tree if it has exactly k leaves. Thus, labeling every vertex with the number of leaves on the subtree rooted at that vertex ensures,

- the root of $\mathcal{T}$ is labeled with k and all other vertices are labeled with positive integers,
- the leaves are labeled with 1, and
- each non-leaf vertex satisfy the property that its label is the sum of the labels of its two children.

Note that, we will think of each non-leaf node with left and right children labeled as a and b as representing the swap graph from $X(a)$ to $X(b)$ for some simplicial complex $X(\leq k)$. Let $\operatorname{Swap}(\mathcal{T}, X)$ be the set of all such swap graphs over X finding representation in the splitting tree $\mathcal{T}$. Indeed the tree $\mathcal{T}$ used in the proof of Theorem VI.3 is just one special instance of a k-splitting tree.

Given a threshold parameter $\tau \leq 1$ and a set of normalized adjacency matrices $\mathcal{A} = \{\mathsf{A}_1, \ldots, \mathsf{A}_s\}$, we define the threshold rank of $\mathcal{A}$ as

$$\operatorname{rank}_{\geq\tau}(\mathcal{A}) := \max_{\mathsf{A}\in\mathcal{A}} \operatorname{rank}_{\geq\tau}(\mathsf{A}),$$

where $\operatorname{rank}_{\geq\tau}(\mathsf{A})$ is denotes usual threshold rank of A as in Definition VII.1.

Now, we are ready to define the notion of a k-CSP instance being $(\mathcal{T}, \tau, r)$-splittable as follows.

Definition VII.2 ($(\mathcal{T}, \tau, r)$-splittability)**.** *A k-CSP instance $\mathfrak{I}$ with the constraint complex $X(\leq k)$ is said to be $(\mathcal{T}, \tau, r)$-splittable if $\mathcal{T}$ is a k-splitting tree and*

$$\operatorname{rank}_{\geq\tau}(\operatorname{Swap}(\mathcal{T}, X)) \leq r.$$

If there exists some k-splitting tree $\mathcal{T}$ such that $\mathfrak{I}$ is $(\mathcal{T}, \tau, r)$-splittable, the instance $\mathfrak{I}$ will be called a (τ, r)-splittable instance.

Now, using this definition we can show that whenever $\operatorname{rank}_{\tau}(\mathfrak{I})$ is bounded for the appropriate choice of τ, after conditioning on a random partial assignment as in Algorithm VI.1 we will have small correlation over the faces $\mathfrak{a} \in X(k)$, i.e.,

Theorem VII.3. *Suppose a simplicial complex $X(\leq k)$ with $X(1) = [n]$ and an L-local PSD ensemble $\mathbf{Y} = \{\mathbf{Y}_1, \ldots, \mathbf{Y}_n\}$ are given. There exists some universal constants $c_4 \geq 0$ and $C'' \geq 0$ satisfying the following: If $L \geq C'' \cdot (q^{4k} \cdot k^7 \cdot r/\varepsilon^5)$, $\operatorname{Supp}(\mathbf{Y}_j) \leq q$ for all $j \in [n]$, and $\mathfrak{I}$ is $(c_4 \cdot (\varepsilon/(4k \cdot q^k))^2, r)$-splittable. Then, we have*

$$\mathop{\mathbb{E}}_{\mathfrak{a}\in X(k)} \left\| \{\mathbf{Y}'_{\mathfrak{a}}\} - \{\mathbf{Y}'_{a_1}\} \cdots \{\mathbf{Y}'_{a_k}\} \right\|_1 \leq \varepsilon, \tag{7}$$

where $\mathbf{Y}'$ is as defined in Algorithm VI.1 on the input of $\{\mathbf{Y}_1, \ldots, \mathbf{Y}_n\}$ and Π_k.

It is important to note that the specific knowledge of the k-splitting tree $\mathcal{T}$ that makes $\mathfrak{I}$ $(\mathcal{T}, \tau, r)$-splittable is only needed for the proof of Theorem VII.3. The conclusion of Theorem VII.3 can be used without the knowledge of the specific k-splitting tree $\mathcal{T}$. Theorem VII.3 can be thought of an extension of Theorem VI.3 to the case where not necessarily every tree is good, and where we can bound the threshold rank instead of the spectral expansion.

This, will readily imply an algorithm

Corollary VII.4. *Suppose $\mathfrak{I}$ is a q-ary k-CSP instance whose constraint complex is $X(\leq k)$. There exists an absolute constant $C'' \geq 0$ and $c_4 \geq 0$ that satisfies the following: If $\mathfrak{I}$ is $(c_4 \cdot (\varepsilon/(4k \cdot q^k))^2, r)$-splittable, then there is an algorithm that runs in time $n^{O\left(\frac{q^{4k}\cdot k^7 \cdot r}{\varepsilon^5}\right)}$ and that is based on $(\frac{C''\cdot k^5 \cdot q^k \cdot r}{\varepsilon^4})$-levels of SoS-hierarchy and Algorithm VI.1 that outputs a random assignment $\eta : [n] \to [q]$ that in expectation ensures $\mathsf{SAT}_{\mathfrak{I}}(\eta) = \mathsf{OPT}(\mathfrak{I}) - \varepsilon$.*

We will need the more general version of Lemma VI.7, proved in [BRS11].

Lemma VII.5 (Lemma 5.4 from [BRS11])**.** [9]*Let $G = (V, E, \Pi_2)$ be a weighted graph, $\{\mathbf{Y}_1, \ldots, \mathbf{Y}_n\}$ a local PSD*

[9]The parameters we get differ from those of [BRS11]. The derivation of our parameters can be found in the appendix of the full version of this paper [AJT19].

ensemble, where we have $\mathrm{Supp}(\mathbf{Y}_i) \le q$ *for every* $i \in V$, *and* $q \ge 0$. *If* $\varepsilon \ge 0$ *is a lower bound on the expected statistical difference between independent and correlated sampling along the edges,i.e.,*

$$\varepsilon \le \mathop{\mathbb{E}}_{\{i,j\}\sim\Pi_2} \|\{\mathbf{Y}_{ij}\} - \{\mathbf{Y}_i\}\{\mathbf{Y}_j\}\|_1.$$

There exists absolute constants $c_3 \ge 0$ *and* $c_4 \ge 0$ *that satisfy the following: Then, conditioning on a random vertex decreases the variances,*

$$\mathop{\mathbb{E}}_{i\sim\Pi_1} \mathop{\mathbb{E}}_{j\sim\Pi_1} \mathop{\mathbb{E}}_{\{\mathbf{Y}_j\}} \mathrm{Var}\left[\mathbf{Y}_i \mid \mathbf{Y}_j\right] \le \mathop{\mathbb{E}}_{i\sim\Pi_1} \left[\mathrm{Var}\left[\mathbf{Y}_i\right]\right] - c_3 \cdot \frac{\varepsilon^4}{q^4 \cdot \mathrm{rank}_{\ge c_4\varepsilon^2/q^2}(G)}.$$

Since we will use this lemma, only with the swap graphs G_{ℓ_1,ℓ_2} and (L/k)-local PSD ensemble $\{\mathbf{Y}_\mathfrak{a}\}_{\mathfrak{a}\in X}$ obtained from the L-local PSD ensemble $\{\mathbf{Y}_1,\ldots,\mathbf{Y}_n\}$, for convenience we will write the corollary we will use more explicitly

Corollary VII.6. *Let* $\ell_1 \ge \ell_2 \ge 0$ *satisfying* $\ell_1 + \ell_2 \le k$ *be given parameters, and let* G_{ℓ_1,ℓ_2} *be the swap graph defined for a* γ*-HDX* $X(\le k)$. *Let* $\{\mathbf{Y}_\mathfrak{a}\}_{\mathfrak{a}\in X}$ *be a local PSD ensemble; and suppose we have* $\mathrm{Supp}(Y_\mathfrak{a}) \le q^k$ *for every* $\mathfrak{a} \in X(\ell_1) \cup X(\ell_2)$ *for some* $q \ge 0$. *Suppose* $\varepsilon > 0$ *satisfies,*

$$\frac{\varepsilon}{4k} \le \mathop{\mathbb{E}}_{\{\mathfrak{s},\mathfrak{t}\}\in E(\ell_1,\ell_2)} \|\{\mathbf{Y}_{\mathfrak{s}\cup\mathfrak{t}}\} - \{\mathbf{Y}_\mathfrak{s}\}\{\mathbf{Y}_\mathfrak{t}\}\|_1.$$

There exists absolute constants $c_3 \ge 0$ *and* $c_5 \ge 0$ *that satisfy the following:*
If $\mathrm{rank}_{\ge c_4\cdot(\varepsilon/(4k\cdot q^k))^2}(G_{\ell_1,\ell_2}) \le r$, *then conditioning on a random face* $\mathfrak{a} \sim \Pi_{\ell_1,\ell_2}$ *decreases the variances, i.e.*

$$\mathop{\mathbb{E}}_{\mathfrak{a},\mathfrak{b}\sim\Pi^2_{\ell_1,\ell_2}} \mathop{\mathbb{E}}_{\{\mathbf{Y}_\mathfrak{a}\}} \mathrm{Var}\left[\mathbf{Y}_\mathfrak{b} \mid \mathbf{Y}_\mathfrak{a}\right] \le \mathop{\mathbb{E}}_{\mathfrak{s}\sim\Pi_{\ell_1,\ell_2}} \left[\mathrm{Var}\left[\mathbf{Y}_\mathfrak{b}\right]\right] - c_5 \cdot \frac{\varepsilon^4}{256\cdot k^4\cdot q^{4k}\cdot r}.$$

ACKNOWLEDGEMENTS

We thank Anand Louis for several enlightening discussions in the initial phases of this work. We are also grateful to the anonymous reviewers for several helpful suggestions. Vedat was supported by NSERC 2950-120715 and 2950-120719 and partially supported by NSF CCF-1254044 and CCF-1718820. Fernando was supported in part by NSF CCF-1254044 and CCF-1816372. Madhur was supported by NSF CCF-1254044 and CCF-1816372.

REFERENCES

[AJQ+19] Vedat Levi Alev, Fernando Granha Jeronimo, Dylan Quintana, Shashank Srivastava, and Madhur Tulsiani. List decoding of direct sum codes. Manuscript, 2019. 3

[AJT19] Vedat Levi Alev, Fernando Granha Jeronimo, and Madhur Tulsiani. Approximating constraint satisfaction problems on high-dimensional expanders. *CoRR*, abs/1907.07833, 2019. URL: http://arxiv.org/abs/1907.07833. 2, 10, 16, 18

[AKK+08] Sanjeev Arora, Subhash Khot, Alexandra Kolla, David Steurer, Madhur Tulsiani, and Nisheeth Vishnoi. Unique games on expanding constraint graphs are easy. In *Proceedings of the 40th ACM Symposium on Theory of Computing*, 2008. 1

[ALGV18] Nima Anari, Kuikui Liu, Shayan Oveis Gharan, and Cynthia Vinzant. Log-concave polynomials II: High-dimensional walks and an FPRAS for counting bases of a matroid. *arXiv preprint arXiv:1811.01816*, 2018. 3

[BH13] Fernando G. S. L. Brandão and Aram Wettroth Harrow. Product-state approximations to quantum ground states. In *Proceedings of the 45th ACM Symposium on Theory of Computing*, pages 871–880, 2013. 3

[BRS11] Boaz Barak, Prasad Raghavendra, and David Steurer. Rounding semidefinite programming hierarchies via global correlation. In *FOCS*, pages 472–481, 2011. 1, 2, 3, 5, 6, 7, 15, 16, 17, 18

[BS02] M. Bernstein and N. J. A. Sloane. Some Canonical Sequences of Integers. *arXiv Mathematics e-prints*, page math/0205301, May 2002. `arXiv:math/0205301`. 9

[BSHR05] Eli Ben-Sasson, Prahladh Harsha, and Sofya Raskhodnikova. Some 3CNF properties are hard to test. *SIAM Journal on Computing*, 35(1):1–21, 2005. 1

[CTZ18] David Conlon, Jonathan Tidor, and Yufei Zhao. Hypergraph expanders of all uniformities from Cayley graphs. *arXiv e-prints*, page arXiv:1809.06342, September 2018. `arXiv:1809.06342`. 2

[DD19] Yotam Dikstein and Irit Dinur. Agreement testing theorems on layered set systems. In *Proceedings of the 60th IEEE Symposium on Foundations of Computer Science*, 2019. URL: https://arxiv.org/abs/1909.00638. 3

[DDFH18] Yotam Dikstein, Irit Dinur, Yuval Filmus, and Prahladh Harsha. Boolean function analysis on high-dimensional expanders. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, pages 38:1–38:20, 2018. 2, 3, 4, 7, 10, 11, 12, 13

[DHK+18] Irit Dinur, Prahladh Harsha, Tali Kaufman, Inbal Livni Navon, and Amnon Ta-Shma. List decoding with double samplers. *Electronic Colloquium on Computational Complexity (ECCC)*, 25:136, 2018. 1, 2

[DK12] Irit Dinur and Tali Kaufman. Locally testable codes and expanders. Manuscript, 2012. 2

[DK17] Irit Dinur and Tali Kaufman. High dimensional expanders imply agreement expanders. In *Proceedings of the 58th IEEE Symposium on Foundations of Computer Science*, pages 974–985, 2017. 1, 2, 4, 7

[Fil16] Yuval Filmus. An orthogonal basis for functions over a slice of the boolean hypercube. *Electr. J. Comb.*, 23(1):P1.23, 2016. 15

[FK96] A. Frieze and R. Kannan. The regularity lemma and approximation schemes for dense problems. In *Proceedings of the 37th IEEE Symposium on Foundations of Computer Science*, 1996. 3

[GM15] Christopher Godsil and Karen Meagher. *Erdős-Ko-Rado Theorems: Algebraic Approaches*. Cambridge Studies in Advanced Mathematics. Cambridge University Press, 2015. 10, 15

[GS11] Venkatesan Guruswami and Ali Kemal Sinop. Lasserre hierarchy, higher eigenvalues, and approximation schemes for graph partitioning and quadratic integer programming with PSD objectives. In *Proceedings of the 52nd IEEE Symposium on Foundations of Computer Science*, pages 482–491, 2011. 1, 3

[GS12] Venkatesan Guruswami and Ali Kemal Sinop. Faster SDP hierarchy solvers for local rounding algorithms. In *Proceedings of the 53rd IEEE Symposium on Foundations of Computer Science*, pages 197–206. IEEE, 2012. 2

[Kho02] Subhash Khot. On the power of unique 2-prover 1-round games. In *Proceedings of the 34th ACM Symposium on Theory of Computing*, pages 767–775, 2002. 1

[KKL16] Tali Kaufman, David Kazhdan, and Alexander Lubotzky. Isoperimetric inequalities for ramanujan complexes and topological expanders. *Geometric and Functional Analysis*, 26(1):250–287, Feb 2016. 2

[KM17] Tali Kaufman and David Mass. High dimensional random walks and colorful expansion. In *Proceedings of the 8th Conference on Innovations in Theoretical Computer Science*, pages 4:1–4:27, 2017. 2

[KM18] Tali Kaufman and David Mass. Good distance lattices from high dimensional expanders. *CoRR*, abs/1803.02849, 2018. URL: http://arxiv.org/abs/1803.02849, arXiv:1803.02849. 1

[KMOW17] Pravesh Kothari, Ryuhei Mori, Ryan O'Donnell, and David Witmer. Sum of squares lower bounds for refuting any CSP. In *Proceedings of the 49th ACM Symposium on Theory of Computing*, 2017. 1

[KO18a] Tali Kaufman and Izhar Oppenheim. Construction of new local spectral high dimensional expanders. In *Proceedings of the 50th ACM Symposium on Theory of Computing*, STOC 2018, pages 773–786. ACM, 2018. 2

[KO18b] Tali Kaufman and Izhar Oppenheim. High order random walks: Beyond spectral gap. In *Approximation, Randomization, and Combinatorial Optimization. Algorithms and Techniques, APPROX/RANDOM 2018, August 20-22, 2018 - Princeton, NJ, USA*, pages 47:1–47:17, 2018. 2, 3, 4

[LLP17] Eyal Lubetzky, Alex Lubotzky, and Ori Parzanchevski. Random walks on Ramanujan complexes and digraphs. *arXiv e-prints*, page arXiv:1702.05452, Feb 2017. arXiv:1702.05452. 2

[LSV05a] Alexander Lubotzky, Beth Samuels, and Uzi Vishne. Explicit constructions of ramanujan complexes of type ad. *Eur. J. Comb.*, 26(6):965–993, August 2005. 2

[LSV05b] Alexander Lubotzky, Beth Samuels, and Uzi Vishne. Ramanujan complexes of typeã. *Israel Journal of Mathematics*, 149(1):267–299, Dec 2005. 2

[MM11] Konstantin Makarychev and Yury Makarychev. How to play unique games on expanders. In *Approximation and Online Algorithms*, pages 190–200. Springer Berlin Heidelberg, 2011. 1

[MR17] Pasin Manurangsi and Prasad Raghavendra. A birthday repetition theorem and complexity of approximating dense csps. In *Proceedings of the 44th International Colloquium on Automata, Languages and Programming*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017. 3

[OGT15] Shayan Oveis Gharan and Luca Trevisan. A new regularity lemma and faster approximation algorithms for low threshold rank graphs. *Theory of Computing*, 11(9):241–256, 2015. URL: http://www.theoryofcomputing.org/articles/v011a009, doi:10.4086/toc.2015.v011a009. 1, 3

[PRT16] Ori Parzanchevski, Ron Rosenthal, and Ran J. Tessler. Isoperimetric inequalities in simplicial complexes. *Combinatorica*, 36(2):195–227, Apr 2016. 2

[RW17] Prasad Raghavendra and Benjamin Weitz. On the bit complexity of sum-of-squares proofs. In *Proceedings of the 44th International Colloquium on Automata, Languages and Programming*. Schloss Dagstuhl-Leibniz-Zentrum fuer Informatik, 2017. 2

[Sag13] B.E. Sagan. *The Symmetric Group: Representations, Combinatorial Algorithms, and Symmetric Functions*. Graduate Texts in Mathematics. Springer New York, 2013. 15

[WJ04] Martin J Wainwright and Michael I Jordan. Treewidth-based conditions for exactness of the Sherali-Adams and Lasserre relaxations. Technical report, Technical Report 671, University of California, Berkeley, 2004. 3

[YZ14] Yuichi Yoshida and Yuan Zhou. Approximation schemes via Sherali-Adams hierarchy for dense constraint satisfaction problems and assignment problems. In *Proceedings of the 5th Conference on Innovations in Theoretical Computer Science*, pages 423–438. ACM, 2014. 3