HIGH DIMENSIONAL HOFFMAN BOUND AND APPLICATIONS IN EXTREMAL COMBINATORICS

YUVAL FILMUS, KONSTANTIN GOLUBEV, AND NOAM LIFSHITZ

Abstract. The n-th tensor power of a graph with vertex set V is the graph on the vertex set V^n, where two vertices are connected by an edge if they are connected in each coordinate. The problem of studying independent sets in tensor powers of graphs is central in combinatorics, and its study involves a beautiful combination of analytical and combinatorial techniques. One powerful method for upper-bounding the largest independent set in a graph is the Hoffman bound, which gives an upper bound on the largest independent set of a graph in terms of its eigenvalues. It is easily seen that the Hoffman bound is sharp on the tensor power of a graph whenever it is sharp for the original graph.

In this paper we introduce the related problem of upper-bounding independent sets in tensor powers of hypergraphs. We show that many of the prominent open problems in extremal combinatorics, such as the Turán problem for (hyper-)graphs, can be encoded as special cases of this problem. We also give a new generalization of the Hoffman bound for hypergraphs which is sharp for the tensor power of a hypergraph whenever it is sharp for the original hypergraph.

As an application of our Hoffman bound, we make progress on the following problem of Frankl from 1990. An extended triangle in a family of sets is a triplet $\{A, B, C\} \subseteq \binom{[n]}{2k}$ such that each element of $[n]$ belongs either to none of the sets $\{A, B, C\}$ or to exactly two of them. Frankl asked how large can a family $\mathcal{F} \subseteq \binom{[n]}{2k}$ be if it does not contain a triangle. We show that if $\frac{1}{2}n \leq 2k \leq \frac{2}{3}n$, then the extremal family is the *star,* i.e. the family of all sets that contains a given element. This covers the entire range in which the star is extremal. As another application, we provide spectral proofs for Mantel's theorem on triangle-free graphs and for Frankl-Tokushige theorem on k-wise intersecting families.

Contents

Date: November 6, 2019.

2010 *Mathematics Subject Classification.* Primary 05C15; Secondary 05C65, 05D05.

Key words and phrases. chromatic number, independence ratio, hypergraph, extremal set theory.

1. Introduction

The celebrated Hoffman bound [23] connects spectral graph theory with extremal combinatorics, by upper-bounding the independence number of a graph in terms of the minimal eigenvalue of its adjacency matrix. The Hoffman bound, in a generalized version due to Lovász [27], has seen many applications in extremal set theory and theoretical computer science.

The Hoffman bound can be used to solve problems in extremal set theory in which the constraints can be modeled as a graph. As an example, the Hoffman bound can be used to prove the fundamental Erdős–Ko–Rado theorem on the size of intersecting families, in which the constraint is that every two sets in the family have nonempty intersection. Other problems involve more complex constraints, and so are not amenable to this method. A simple example is the s-wise intersecting Erdős–Ko–Rado theorem, due to Frankl [15], which concerns families in which every s sets have nonempty intersection. In this case the constraints can be modeled as a *hypergraph* rather than as a graph.

Recently, Hoffman's bound has been generalized to hypergraphs [4, 20]. The new bound is particularly attractive for upper-bounding independent sets in tensor powers of hypergraphs, a setting which we describe in detail below. We demonstrate the power of this method by solving a problem of Frankl on triangle-free families and by giving a spectral proof of Mantel's theorem. We also formulate a number of known problems in the language of independent sets in hypergraphs.

1.1. **Notations.**

A multiset is an unordered collection of elements that is allowed to have repetitions, its size is the number of its elements counting the multiplicity. An *i-multiset* is a multiset of size i. Let V be a set. We denote by $V^{[i]}$ the collection of all i-multisets of elements of V, and elements of $V^{[i]}$ will be denoted by $[v_1, \ldots, v_i]$. $V^{[0]}$ consists of the the empty set.

Definition 1.1. A *weighted k-uniform hypergraph* is a pair $X = (V, \mu)$ where V is the vertex set and μ is a probability distribution on $V^{[k]}$.

For $0 \le i \le k-1$, define a probability measure μ_i on $V^{[i]}$ by the following process. First, choose a multiset $[v_1, \ldots, v_k]$ according to μ, and then choose an

i-submultiset of it uniformly at random. We write $X^{(i)}$ for the set of elements of $V^{[i]}$ whose μ_i measure is positive. The elements of $X^{(i)}$ are called the *i-faces* of X, and the elements of $X^{(0)} \cup \cdots \cup X^{(k)}$ are called the *faces* of X. Note that if σ_2 is a face of X, then σ_1 is a face of X for any $\sigma_1 \subseteq \sigma_2$. Note that $X^{(0)} = \{\emptyset\}$, i.e., the empty set is the one and only 0-face of X. Therefore, the collection of multisets $\left(X^{(k)}, X^{(k-1)}, \ldots, X^{(0)}\right)$ can be viewed as an abstract simplicial complex of dimension $(k-1)$ (e.g., as defined in [4,20]) which is however allowed to have loops (multiples of a vertex in a face). Conversely, an abstract simplicial complex can be made into a weighted uniform hypergraph by introducing a probability measure which is positive on its maximal faces.

Definition 1.2. A set $I \subseteq V$ is said to be *independent in a k-uniform hypergraph X* if no k-face of X is contained in I. The largest possible value of $\mu_0(I)$, where $I \subseteq V$ is an independent set in X, is called the *independence number* of X and denoted $\alpha(X)$. A subset $I \subseteq V$ is said to be an *extremal independent set* of X if $\mu_0(I) = \alpha(X)$.

We can couple the distributions $\mu_0, \ldots, \mu_k$ into a distribution $\boldsymbol{\mu}$ over flags of X by sampling a random k-face according to μ and removing elements from it one by one uniformly at random. It will also be useful to consider distributions $\tilde{\mu}_i$ on $V^{[i]}$, obtained by sampling an i-face according to μ_i, and then choosing a random order of the vertices it contains.

Let $\sigma \in X^{(i)}$ be an i-face. Its *link* in X is the $(k-i)$-uniform hypergraph $X_\sigma = (V, \mu_\sigma)$, where μ_σ is the probability distribution that corresponds to the following process: sample a random flag according to $\boldsymbol{\mu}$ subject to $\boldsymbol{\mu}_i = \sigma$, and output $\boldsymbol{\mu}_k \setminus \sigma$. For a set $A \in X^{(k-i)}$ we shall say that $\mu_\sigma(A)$ is the *relative measure of A according to σ*. Note that the link of the empty set is the whole hypergraph X itself.

The *skeleton* of X is the weighted graph $S(X)$ on the vertex set $X^{(1)}$, whose edges are $X^{(2)}$, and whose weights are given by μ_2. The inner product on the space $L^2\left(X^{(1)}, \mu_1\right)$ of functions on the vertices is defined as

$$\langle f, g\rangle = \mathbb{E}_{v \sim \mu_1} f(v) g(v) = \sum_{v \in X^{(1)}} f(v) g(v) \mu_1(v).$$

The normalized adjacency operator T_X of X is that of the skeleton $S(X)$. In other words, T_X acts on $L^2\left(X^{(1)}, \mu_1\right)$ as follows:

$$\left(T_X f\right)(v) = \mathbb{E}_{\mu_1(X_v)}[f].$$

If $f, g \in L^2(X^{(1)}, \mu_1)$ then

$$\begin{aligned} \langle f, T_X g\rangle &= \mathbb{E}_{\boldsymbol{v} \sim \mu_1} f(\boldsymbol{v}) \mathbb{E}_{\boldsymbol{u} \sim \mu_1(X_{\boldsymbol{v}})} g(\boldsymbol{u}) \\ &= \mathbb{E}_{\boldsymbol{w} \sim \boldsymbol{\mu}} f(\boldsymbol{w}_1) g(\boldsymbol{w}_2 \setminus \boldsymbol{w}_1) = \mathbb{E}_{(\boldsymbol{u}, \boldsymbol{v}) \sim \tilde{\mu}_2} f(\boldsymbol{v}) g(\boldsymbol{u}). \end{aligned}$$

This shows that T_X is self-adjoint, and so has real eigenvalues. The matrix form of T_X is given by the formula

$$T_X(u, v) = \begin{cases} \frac{\mu_2([u,u])}{\mu_1(u)} & \text{if } u = v; \\ \frac{\mu_2([u,v])}{2\mu_1(u)} & \text{if } u \neq v. \end{cases} \tag{1.1}$$

Similar reasoning shows that we can sample $[u, v] \sim \mu_2$ by sampling $u \sim \mu_1$ and $v \sim \mu_1(X_v)$.

Note that if V is a finite set (which is the case throughout this paper), then the largest eigenvalue of T_X is 1 and is achieved on the constant function. By $\lambda(X)$ we denote the smallest eigenvalue of T_X. For all $0 \le i < k-2$, we write

$$\lambda_i(X) = \min_{\sigma \in X^{(i)}} \left[\lambda\left(S\left(X_\sigma\right)\right)\right].$$

In other words, $\lambda_i(X)$ is the minimal possible value of an eigenvalue of the normalized adjacency matrix of a skeleton of the link of an i-face of X. Note that $\lambda_0(X)$ is just the smallest eigenvalue of the normalized adjacency operator on the skeleton of X.

Definition 1.3. The tensor product $X \otimes X'$ of two k-uniform hypergraphs $X = (V, \mu)$ and $X = (V', \mu')$ is a k-uniform hypergraph $(V \times V', \mu \times \mu')$, where $\mu \times \mu'$ stands for the product measure on $(V \times V')^{[k]} \simeq V^{[k]} \times V'^{[k]}$. For a k-uniform hypergraph X, we denote by $X^{\otimes n} = \underbrace{X \otimes \cdots \otimes X}_{n}$ its n-th tensor power.

1.2. **Results.** We prove a new upper bound for the independence number of a hypergraph, and its invariance under the tensor power operation.

Theorem 1.4. *Let $X = (V, \mu)$ be a k-uniform hypergraph. Then*

$$\alpha(X) \le 1 - \frac{1}{(1-\lambda_0)(1-\lambda_1)\cdots(1-\lambda_{k-2})}. \tag{1.2}$$

If in addition $\lambda_i \le 0$ for all $0 \le i < k-2$, then for any positive integer n the following inequality holds for $X^{\otimes n}$:

$$\alpha(X^{\otimes n}) \le 1 - \frac{1}{(1-\lambda_0)(1-\lambda_1)\cdots(1-\lambda_{d-1})}.$$

In particular, if the bound is sharp for X, it remains sharp for its tensor powers, as $\alpha(X^{\otimes n}) = \alpha(X)$.

We apply Theorem 1.4 to deduce the following result on the Frankl's problem on triangle-free families, in both a uniform and a p-biased versions.

Theorem 1.5. ***The uniform version.*** *Let $\binom{[n]}{2k}$ be the space of $2k$-subsets of $[n]$, where $n \le 4k-1$. If $\mathcal{F} \subseteq \binom{[n]}{2k}$ is a family of subsets which does not contain three distinct subsets whose symmetric difference is empty, then $|\mathcal{F}| \le \binom{n-1}{2k-1}$. This bound is sharp, as, for example, the family of all subsets containing the element 1 satisfies the condition and contains $\binom{n-1}{2k-1}$ subsets.*

The p-biased version. *Let $\{0,1\}^n$ be the space of $\{0,1\}$-vectors of length n endowed with be the p-biased measure μ, where $1/2 \le p \le 2/3$. If $\mathcal{F} \subseteq \{0,1\}^n$ is a family of vectors which does not contain three distinct vectors whose sum is zero, then $\mu(\mathcal{F}) \le p$. This bound is sharp, as, for example, the set of all vectors having 1 as their first coordinate satisfies the condition and has measure p.*

Our method also provides spectral proves of Mantel's theorem on triangle-free graphs and Frankl-Tokushige theorem on k-wise intersecting families.

Theorem 1.6. *[29] If a graph on n vertices contains no triangle, then it contains at most $\left\lfloor \frac{n^2}{4} \right\rfloor$ edges.*

Theorem 1.7. *[16] Let $k \geq 2$ and $p \leq 1 - \frac{1}{k}$. Assume $\mathcal{F} \subset \mathcal{P}([n])$ is k-wise intersecting family of subsets of $[n]$, that is, for all $F_1, \ldots, F_k \in \mathcal{F}$*

$$F_1 \cap \cdots \cap F_k \neq \emptyset.$$

Then $\mu_p(\mathcal{F}) = \sum_{F \in \mathcal{F}} p^{|F|}(1-p)^{n-|F|} \leq p$, in other words, μ_p is the p-biased measure on $\mathcal{P}([n])$.

1.3. **Structure of the Paper.** In Section 2, we give a brief overview of the method for graphs: the Hoffman bound, its behavior for tensor product of graphs, and applications in extremal combinatorics. In Section 3, we introduce the required hypergraph definitions and notations, and translate a number of known problems to the language of independent sets of hypergraphs. In Section 4, we prove Theorem 1.4 and compare the new bound to the known ones. In Sections 5 and 6, we prove Theorems 1.5 and 1.6, respectively.

2. Hoffman bound for graphs

Let $G = (V, \mu)$ be a graph, that is, a 2-uniform hypergraph in the notations of this paper. By edges of G we mean the support of μ in $V^{[2]}$. A subset of V is called *independent* if it does not contain any edges. Recall that μ_1 is the induced measure on V. The *independence number* of G is the maximal value of $\mu_1(A)$ over all independent sets $A \subseteq V$. If $\mu_1(A) = \alpha(G)$, we say that A is an *extremal independent set* of G. We write $\lambda_{\min}(G)$ for the minimal eigenvalue of the normalized adjacency T_G on G, see Equation 1.1 for the formula.

The Hoffman bound for graphs gives an upper bound on the largest independent set of G in terms of its minimal eigenvalue.

Theorem 2.1 (Hoffman bound, [23])**.** *Let G be a graph. Then*

$$\alpha(G) \leq \frac{-\lambda_{\min}(G)}{1 - \lambda_{\min}(G)}.$$

Ever since Hoffman's original work, the Hoffman bound has become a central tool in the study of intersection problems in combinatorics. For example, Lovász, [27] showed that it can be used to prove the celebrated Erdős–Ko–Rado theorem, [13], (aka EKR theorem), which states that for $n \geq 2k$, the maximum size of an intersecting family of k-element subsets of $[n]$ is $\binom{n-1}{k-1}$. The theorem follows directly from the Hoffman bound by considering the Kneser graph, which is the graph on the vertex set $\binom{[n]}{k}$ in which two sets are connected if they are disjoint. It is crucial for the proof that the Hoffman bound is tight in this case.

The Hoffman bound can be also used in a more sophisticated manner. For instance, in [36], Wilson considered the t-intersection variant of the EKR theorem, in which the goal is to find the largest subfamily of $\binom{[n]}{k}$ in which any two sets have at least t elements in common. In contrast to the EKR theorem, in this case applying Hoffman's bound on the generalized Kneser graph (in which two sets are connected if their intersection contains less than t points) does not give a tight upper bound. Instead, one needs to accurately choose weights for the edges of this graph, some of them negative, and only then apply the Hoffman bound. In this way, Wilson managed to determine all values of n, k, t in which the extremal family is the family of all sets that contain a given set of size t, namely, $n \geq (t+1)(k-t+1)$.

This approach became more systematic in the work of Friedgut, [17], who applied Fourier analysis to construct a matrix for the p-biased version of Wilson theorem.

This Fourier-analytic approach was also useful in the proof of a long-standing problem of Simonovitz and Sós on triangle-intersecting families of graphs, [10]. Finally, Ellis, Friedgut, and Pilpel, [11], used a similar approach to solve an old problem of Deza and Frankl, showing that a t-intersecting family of permutations in S_n contains at most $(n-t)!$ permutations, for large enough n (two permutations t-intersect if they agree on the image of at least t points).

The main flaw of the Hoffman bound is the fact that it does not apply to problems in which the constraint involves more than two sets. For example, it cannot be used to upper-bound the size of a subfamily of $\binom{[n]}{k}$ in which any *three* sets t-intersect. It is therefore of great importance to find high-dimensional analogues of the Hoffman bound that fit these more general restrictions.

Over the years there has been great interest in the problem of generalizing results from graphs to hypergraphs (or, equivalently, simplical complexes), see e.g. [21, 22, 26, 28, 32, 33, 34]. In particular, two generalizations of the Hoffman bound were given by [4] and [20]. Such results are known as *high-dimensional Hoffman bounds.* The goal of this paper is to give a new high-dimensional Hoffman bound, which we believe is the right tool for tackling many problems in extremal combinatorics.

2.1. **Independent sets in tensor power of graphs.** The Hoffman bound is particularly useful for independent sets in *tensor powers of graphs.* Given graphs $G = (V, \mu)$, $G' = (V', \mu')$, their tensor product is the graph $G \otimes G'$ whose vertex set is $V \times V'$ endowed with the product measure $\mu \times \mu'$. In particular, two of vertices in the product are connected by an edge if they are connected by an edge in each coordinate. The n-th tensor power of G is the graph $G^{\otimes n} = G \otimes \cdots \otimes G$. Independent sets in tensor products of graphs are well studied, see e.g. [2], [3], and [7].One motivation is their connection to the theory of hardness of approximation, see [8]. However, our main focus in this paper will be the connection to extremal set theory. This connection was first implicitly established by Friedgut [18], and later more explicitly by Dinur and Friedgut [6].

It is a well-known rule of thumb (backed by various results) that the two Erdős–Rényi models of random graphs, $G(n, p)$ and $G(n, m)$, should behave similarly when $p = \frac{m}{n}$. A similar phenomenon holds in extremal set theory: questions about subfamilies of $\binom{[n]}{k}$ behave similarly to questions about subsets of $\mathcal{P}([n])$ with respect to the p-biased measure μ_p for $p = k/n$, where a set $\boldsymbol{A} \subseteq [n]$ is chosen by independently putting each element of $[n]$ inside $\boldsymbol{A}$ with probability p and outside of it with probability $1 - p$. We write $\mu_p(\mathcal{F})$ for the probability that a random subset $\boldsymbol{A} \sim \mu_p$ belongs to $\mathcal{F}$.

For instance, the EKR theorem, which asks how large can an intersecting family $\mathcal{F} \subseteq \binom{[n]}{k}$ be, can be rephrased as follows:

> Suppose that $\mathcal{F}$ is intersecting. How large can the probability that a random k-set belongs to $\mathcal{F}$ be?

This formulation of the problem immediately suggests the following p-biased analogue:

> How large can $\mu_p(\mathcal{F})$ be if $\mathcal{F} \subseteq \mathcal{P}([n])$ is intersecting?

This problem was first studied by Ahlswede and Katona, [1], in the 70's. While both the EKR problem and its p-biased analogue have already been solved, this connection is still useful for many of the generalizations of the EKR theorem, where a result in the k-uniform setting can be deduced from the p-biased setting and vice versa, see [9].

The p-biased version of the EKR theorem is the problem of determining the extremal independent sets in the graph $G^{\otimes n}$, where G consists of two vertices $\{0,1\}$ with the (undirected) edge $\{1,0\}$ having the weight $2p$, and with the edge $\{0,0\}$ having the weight $1-2p$. This observation was used by Friedgut, [17], Dinur and Friedgut, [9], and later by Friedgut and Regev, [19], to apply Fourier-analytic methods that are natural in the context of graph products in order to study variants of the EKR theorem.

In a suitable basis, the matrix of the adjacency operator of a tensor product of graphs is the Kronecker product of the adjacency operator matrices of the factors. Hence, the following property holds

$$\lambda_{\min}\left(G^{\otimes n}\right)=\begin{cases}\lambda_{\min}\left(G\right) & \text{if } \lambda_{\min}\left(G\right)\le 0;\\ \lambda_{\min}\left(G\right)^{n} & \text{if } \lambda_{\min}\left(G\right)\ge 0.\end{cases}$$

This immediately implies that the Hoffman bound is sharp on $G^{\otimes n}$, whenever it is sharp on G (the Hoffman bound is never sharp if $\lambda_{\min}$ is positive). This reduces the p-biased EKR problem for families $\mathcal{F}\subseteq\mathcal{P}\left([n]\right)$ to the problem of showing that the Hoffman bound is sharp in the special case $n=1$, which can be verified directly.

A subset A of $G^{\otimes n}$ is called a *dictatorship* if there exists a set $B\subseteq G$ and $1\le i\le n$ such that a vertex $x=(x_1,\ldots,x_n)$ is in A iff x_i is in B. The above observation shows that if the Hoffman bound is sharp for G, then the Hoffman bound is sharp for $G^{\otimes n}$ as well, and the dictatorships corresponding to extremal independent sets of G are extremal for $G^{\otimes n}$ (not necessarily exclusively). Alon, Dinur, Friedgut, and Sudakov, [2], showed the following stronger version of this observation.

Theorem 2.2 ([2])**.** *Let G be a weighted connected non-bipartite graph. If the Hoffman bound is sharp for G, i.e., $\alpha\left(G\right)=\frac{-\lambda_{\min}}{1-\lambda_{\min}}$, then $\alpha\left(G^{\otimes n}\right)=\alpha\left(G\right)$. Moreover, if A is an independent set with $\mu_G\left(A\right)=\alpha\left(G\right)$, then A is a dictatorship.*

Suppose additionally that $\mu_G\left(\{v\}\right)=\Theta\left(1\right)$ for each $v\in G$. Then for each $\epsilon>0$ there exists $\delta>0$ such that if an independent set A satisfies $\mu_G\left(A\right)>\alpha\left(G\right)-\delta$, then there exists an independent dictatorship B such that $\mu_G\left(A\Delta B\right)<\epsilon$.

3. Known problems in the language of hypergraphs

There are plenty of reasons why the above theory begs to be generalized to hypergraphs (or, equivalently, simplical complexes). In addition to the definitions given in the introduction, let us give a version of them in the k-partite setting.

Definition 3.1. A weighted k-partite hypergraph is a tuple $X=(V_1,\ldots,V_k,\mu)$, where μ is a probability distribution on $V_1\times\cdots\times V_k$. The probability distribution μ_{X,V_i} is the probability distribution on V_i, where a vertex is chosen as the projection on V_i of a random element chosen according to μ. Sets $A_1\subseteq V_1,\ldots,A_k\subseteq V_k$ are said to be *cross-independent* if the probability that a random element of μ belongs to $\prod_{i=1}^{k}A_i$ is 0. The tensor power of X is the k-partite hypergraph $X^{\otimes n}=\left(V_1^n,\ldots,V_k^n,\mu^{\otimes n}\right)$, where $\mu^{\otimes n}$ is the product probability distribution.

Independent sets in tensor powers of hypergraphs arise all over combinatorics, and many of the fundamental problems in extremal combinatorics can be formulated as problems about independent sets in tensor powers of hypergraphs. Below we formulate several well-known problems, solved and open, in the language of hypergraphs, and give proof to two of them: Frankl's Triangle Problem and Mantel's Theorem.

3.1. **Number theory.** How large can a subset $A \subseteq \mathbb{F}_q^n$ be if it does not contain elements $x_1, \ldots, x_m \in A$ that form a solution to the system of equations $\sum_{i=1}^m a_{ij} x_i = b_j$, for parameters $a_{ij}, b_j \in \mathbb{F}_q$? For instance, the Meshulam–Roth theorem, [30], concerns the problem of determining how large can a subset of $\mathbb{F}_p^n$ be if it contains no non-trivial solutions to the equation $x + z = 2y$. Similarly, the analogous problem for k-term arithmetic progressions can be formulated as a non-trivial solution to the equations

$$x_1 + x_3 = 2x_2, x_2 + x_4 = 2x_3, \ldots, x_{m-2} + x_m = 2x_{m-1}.$$

Take the hypergraph $X = (\mathbb{F}_q, \mu)$, where μ is positive on the solutions to the equations $\sum_{i=1}^m a_{ij} x_i = b_j$. Since a solution to the system of equations in $\mathbb{F}_q^n$ is a solution iff it is a solution in each coordinate, the hypergraph $X^{\otimes n}$ is the hypergraph whose independent sets correspond to solutions of the same system of equations as above. This observation was first given by Mossel, [31].

3.2. **Turán problem for hypergraphs.** One of the fundamental problems in extremal combinatorics is the Turán problem for hypergraphs, which asks how large can a family $\mathcal{F} \subseteq \binom{[n]}{k}$ be if it does not contain a copy of some given hypergraph H. The Turán problem can be restated as a problem about independent sets in tensor powers of k-partite-hypergraphs, see [25]. Let us state it in the case of triangles in graphs.

Example 3.2. Let $K_{2,2,2}$ be the complete 3-partite hypergraph between sets A_1, A_2, A_3 of size 2, let E_{ij} be the set of edges between A_i and A_j, and let $X = (V_1, V_2, V_3, \mu)$ be the 3-partite hypergraph, where

$$V_1 = E_{23}, V_2 = E_{31}, V_3 = E_{12},$$

and where μ is the uniform measure on the set of triangles in $K_{2,2,2}$. The hypergraph

$$X^{\otimes n}$$

is the hypergraph whose vertices correspond to the edges in $K_{2^n,2^n,2^n}$ and whose edges correspond to the triangles in $K_{2^n,2^n,2^n}$. Therefore, cross-independent sets correspond to three directed graphs

$$G_1 \subseteq V_1 \times V_2, G_2 \subseteq V_2 \times V_3, G_3 \subseteq V_3 \times V_1,$$

where each V_i is $\{0,1\}^n$.

As mentioned above, this example can be generalized to arbitrary hypergraphs. In Section 6, we shall use this construction to give a spectral proof of Mantel's theorem for graphs with 2^n vertices. In this context, Mantel's theorem can be restated as follows:

Theorem 3.3. *Let G_1, G_2, G_3 be cross-independent sets in $X^{\otimes n}$. Suppose additionally that G_1, G_2, G_3 are all equal to some bipartite graph $G \subseteq (\{0,1\} \times \{0,1\})^n$, and that G corresponds to a graph in the sense that $(a,b) \in G$ if and only if $(b,a) \in G$ (in other words, G is the bipartite cover of some graph). Then the largest value of $|G|$ is attained when G is the dictatorship of all $x \in (\{0,1\} \times \{0,1\})^n$ whose first coordinate is either $(0,1)$ or $(1,0)$.*

Indeed, the dictatorship of all $x \in (\{0,1\} \times \{0,1\})^n$ whose first coordinate is either $(0,1)$ or $(1,0)$ corresponds to a balanced complete bipartite graph, which is extremal for Mantel's theorem. Interestingly enough, the dictatorships contain only some of the complete balanced bipartite graphs, but not all of them.

3.3. **Extremal set theory.** Many p-biased versions of problems in extremal set theory can be described as a special case of the problem: "How large can $\alpha\left(X^{\otimes n}\right)$ be?", where X is a weighted hypergraph.

3.3.1. *Erdős Matching Conjecture.* Our first example is the Erdős Matching Conjecture, [12], from 1965. An *s-matching* is a family of s sets $\{A_1, \ldots, A_s\}$ that are pairwise disjoint. Erdős asked how large can a family $\mathcal{F} \subseteq \binom{[n]}{k}$ be if it does not contain an s-matching. He conjectured that the extremal family is either the family of all sets that intersect a given set of size $s-1$ in at least one element, or the family $\binom{[ks-1]}{k}$. The corresponding p-biased version of this problem is as follows:

> Given $p \leq \frac{1}{s}$, how large can $\mu_p(\mathcal{F})$ be if $\mathcal{F}$ does not contain an s-matching?

This is the problem of determining the independence number of the n-th tensor power of the s-uniform hypergraph whose vertex set is $\{0,1\}$ with the weight function $\mu([1,0,\ldots,0]) = sp$ and $\mu([0,\ldots,0]) = 1-sp$ (recall that $[\cdot]$ is our notation for multiset). A nice feature of the p-biased variant of the Erdős Matching Conjecture is that there is only one suggestion for the extremal family, which is the family of all sets that intersect a given set of size $s-1$ in at least one element.

3.3.2. *s-wise Intersecting Families.* The second example is the problem of s-wise intersecting families, first studied by Frankl, [15]. A family $\mathcal{F} \subseteq \binom{[n]}{k}$ is *s-wise intersecting* if the intersection of every s sets in $\mathcal{F}$ is nonempty. Frankl showed that when $k \leq \frac{s-1}{s}n$, the extremal s-wise intersecting family is the family of all sets that contain a given element (otherwise every family is s-wise intersecting). The p-biased version of the problem was studied by Frankl and Tokushige, [16]. They showed that the largest value of $\mu_p(\mathcal{F})$ for an s-wise intersecting family $\mathcal{F}$ is p, as long as $p \leq \frac{s-1}{s}$. This problem can be expressed as the determining the independence number of the hypergraph $X^{\otimes n}$, where

(1) The hypergraph $X = (V, \mu)$ has $\{0,1\}$ as its vertex set V.
(2) The induced distribution μ_1 on V is the p-biased one.
(3) $\mu(x) = 0$ if and only x is the all ones vector.

It is easy to construct many hypergraphs X that satisfy these hypotheses. We reprove this result in 4.2.

3.3.3. *Frankl's Turán Problem.* Last but not least, this problem is related to Frankl's Turán problem on hypergraphs without extended triangles. A *triangle* in $\mathcal{P}([n])$ is a $2k$-uniform hypergraph $\{A, B, C\}$ such that each element of $[n]$ belongs to an even number of the sets A, B, C. In other words, there exist disjoint k-element sets D, E, F such that $D \cup E = A$, $D \cup F = B$, and $E \cup F = C$. Frankl, [14], asked how large can a family $\mathcal{F} \subseteq \binom{[n]}{2k}$ be if it does not contain a triangle. The reason for considering only even uniformities is that no k-uniform triangle exists for an odd k. The p-biased version of the problem is as follows:

> Given $p \leq \frac{2}{3}$, how large can $\mu_p(\mathcal{F})$ be if $\mathcal{F} \subseteq \mathcal{P}([n])$ does not contain a triangle?

The reason for the condition $p \leq \frac{2}{3}$ is the fact that the family $\left\{A : |A| > \frac{2}{3}n\right\}$ is triangle-free, and its p-biased measure tends to 1 as n tends to infinity.

The p-biased version of Frankl's Turán problem is the problem of determining the independence number of the hypergraph $X^{\otimes n}$, where $X = (V, \mu)$ is with $V = \{0, 1\}$ and

$$\mu([1,1,0]) = \frac{3}{2}p, \ \mu([0,0,0]) = 1 - \frac{3}{2}p.$$

In Section 5 we prove the following theorem:

Theorem 3.4. *If a family of $2k$-subsets of $[n]$ contains no three distinct subsets whose symmetric difference is empty, then the family contains at most $\frac{2k}{\min(n,4k-1)}\binom{n}{2k}$ subsets.*

Furthermore, when $n \leq 4k-1$ and $p \geq 1/2$ the bounds are tight for "dictatorships" (all subsets or vectors containing a specific point), and otherwise the bounds are asymptotically tight, in the p-biased case for the family of all vectors having odd parity, and in the $2k$-uniform case for the family of all subsets whose intersection with $[\lfloor n/2 \rfloor]$ is odd.

Similarly we prove that if a subset of $\{0,1\}^n$ contains no three distinct vectors summing to zero, then for all $p \leq 2/3$, its μ_p-measure is at most $\max(p, 1/2)$.

4. High dimensional Hoffman bound

Suppose we are given a problem in extremal set theory where constraints on more than two elements are involved. A possible strategy for solving it is to first incorporate families $\mathcal{F}$ that satisfy the constraint as independent sets in some hypergraph or simplical complex, and then to find and apply a high-dimensional generalization of the Hoffman bound in order to bound the size of $\mathcal{F}$. Two such generalizations of the Hoffman bound were obtained recently by Golubev, [20], and by Bachoc, Gundert, and Passuello, [4]. However, none of them seem to give sharp results in our problems of interest. Instead, we develop a different generalization of the Hoffman bound in the spirit of [20].

Let $X = (V, \mu)$ be a k-uniform hypergraph on the vertex set V. Recall (see Subsection 1.1) that for $0 \leq i \leq k-2$ we denote by $\lambda_i(X)$ the minimal possible value of an eigenvalue of the normalized adjacency matrix of the skeleton of the link of an i-face of X. That is,

$$\lambda_i(X) = \min_{\sigma \in X^{(i)}} \left[\lambda\left(S\left(X_\sigma\right)\right)\right].$$

Note that the hypergraph itself is the link of the only 0-face, the empty set, and hence $\lambda_0(X)$ is just the smallest eigenvalue the normalized adjacency operator on the skeleton of X.

Example. Let $X = (\{1,2\}, \mu)$ be a graph (in other words, 2-uniform hypergraph) on two vertices $\{1,2\}$ with the probability measure μ defined on the edges as

$$\mu([1,1]) = p_1, \ \mu([1,2]) = p_2, \ \mu([2,2]) = p_3.$$

Then the induced distribution μ_1 on the vertices is as follows:

$$\mu_1(1) = p_1 + \frac{1}{2}p_2, \ \mu_1(2) = \frac{1}{2}p_2 + p_3.$$

The normalized adjacency operator T_X on the skeleton of X has the matrix form

$$T_X = \begin{pmatrix} \frac{p_1}{p_1+\frac{1}{2}p_2} & \frac{\frac{1}{2}p_2}{p_1+\frac{1}{2}p_2} \\ \frac{\frac{1}{2}p_2}{\frac{1}{2}p_2+p_3} & \frac{p_3}{\frac{1}{2}p_2+p_3} \end{pmatrix}$$

and while its largest eigenvalue is equal to 1, the smallest one is equal to

$$\lambda_0(X) = 1 - \frac{2p_2}{1-(p_1-p_3)^2}.$$

Theorem 4.1. *Let $X = (V, \mu)$ be a k-uniform hypergraph. Then*

$$\alpha\left(X\right) \le 1 - \frac{1}{\left(1-\lambda_0\right)\left(1-\lambda_1\right)\cdots\left(1-\lambda_{k-2}\right)}. \tag{4.1}$$

Proof. The proof goes by induction on the uniformity of the hypergraph. The base case, $k = 2$, is the graph case, and the bound 4.1 reads as the classical Hoffman bound. Assume that the bound holds for $(k-1)$-uniform hypergraphs. Let T_X be the normalized adjacency operator of the skeleton of X, and let $v_1 = 1, v_2, \ldots, v_m$ an orthonormal basis of its eigenvectors with eigenvalues $1 \ge l_2 \ge \cdots \ge l_m = \lambda_0$ (recall that T_X is self-adjoint). Let I be an independent set of measure $\alpha(X)$, and let $f = 1_I$ be its indicator function. We may write

$$f = \sum_{i=1}^{m} \langle f, v_i \rangle\, v_i.$$

On the one hand,

$$\langle T_X f, f\rangle = \Pr_{[\boldsymbol{x},\boldsymbol{y}]\sim\mu_2}\left[\boldsymbol{x},\boldsymbol{y} \in I\right],$$

or in other words, it is equal to the probability of an ordered edge (1-face) distributed according to μ_2 to have both ends in I. On the other hand,

$$\begin{aligned}\langle T_X f, f\rangle &= \sum_{i=1}^{m} l_i \left\langle f, v_i\right\rangle^2 \\ &\ge \left\langle f, 1\right\rangle^2 + \sum_{i=2}^{m} \lambda_0 \left\langle f, v_i\right\rangle^2 \\ &= \left\langle f, 1\right\rangle^2 (1-\lambda_0) + \lambda_0 \left\langle f, f\right\rangle \\ &= \mathbb{E}\left[f\right]^2 (1-\lambda_0) + \lambda_0 \mathbb{E}\left[f^2\right].\end{aligned}$$

Since f is an indicator function, $\mathbb{E}[f] = \mathbb{E}[f^2] = \alpha(X)$, and so

$$\Pr_{[\boldsymbol{x},\boldsymbol{y}]\sim\mu_2}\left[\boldsymbol{x},\boldsymbol{y} \in I\right] \ge \alpha(X)^2 (1-\lambda_0) + \lambda_0 \alpha(X).$$

Note that

$$\Pr_{[\boldsymbol{x},\boldsymbol{y}]\sim\mu_2}\left[\boldsymbol{x},\boldsymbol{y} \in I\right] \le \alpha(X) \cdot \max_{x\in I} \Pr_{\boldsymbol{y}\sim\mu_1(X_x)}\left[\boldsymbol{y} \in I\right]$$

and that for a fixed vertex x, the probability $\Pr_{y\sim\mu_1(X_x)}\left[\boldsymbol{y} \in I\right]$ is the measure of an independent set of vertices in its link X_x, which is a $(k-1)$-uniform hypergraph. By the induction assumption,

$$\Pr_{y\sim\mu_1(X_x)}\left[\boldsymbol{y} \in I\right] \le 1 - \frac{1}{\left(1-\lambda_1\right)\cdots\left(1-\lambda_{k-2}\right)}.$$

Combining the above bounds, we deduce that

$$(1-\lambda_0)\alpha(X)+\lambda_0 \leq \Pr_{y\sim\mu_1(X_x)}[\boldsymbol{y}\in I] \leq 1-\frac{1}{(1-\lambda_1)\cdots(1-\lambda_{k-2})},$$

which proves the required bound after rearrangement. □

4.1. **Comparison to prior work.** There are two known spectral upper bounds on the independence number of a hypergraph (equivalently, a simplicial complex): one in [20], to which we will refer as *the Laplacian bound*, and one in [4], to which we will refer as *the Theta bound*. To the bound in Theorem 4.1 we will refer as *the Link bound*.

The Link bound and the Laplacian bound are based on the same idea. Namely, the bound on the size of an independent set is obtained via a combination of a lower and an upper bound on the number of 2-faces between the maximal independent set and its complement. Below we provide an explanation why the Link bound is always as least as good as the Laplacian bound.

The Theta bound follows a different approach. In [4], the authors show that the Theta bound is incomparable to the Laplacian bound by providing two families of simplicial complexes, on which one bound is sharp while the other is not, and vice versa. However, the new bound is sharp for all examples provided in [4]. In other words, the Link bound is better than the Theta bound in at least one case, and they are equal in other cases. It is not clear whether these bounds are comparable or not.

In order to compare the Laplacian bound with the Link one, let us reformulate the former in the terms of the normalized Laplacian (see, e.g., [24] for the definition of the normalized Laplacian). Let X be a k-uniform hypergraph. Then

$$\alpha(X) \leq 1-\frac{1}{\mu_0\cdots\mu_{k-2}},$$

where μ_i is the largest eigenvalue of the normalized i-Laplacian on X. In order to show that the Link bound is at least as good as the Laplacian one, it suffices to show that $(1-\lambda_i)\leq\mu_i$ for all $i=0,\ldots,k-2$. For graphs, the normalized Laplace operator Δ and the normalized adjacency operator T satisfy $\Delta = Id - T$, and hence $\lambda_0 = 1-\mu_0$. For $0<i$, the Laplace bound exploits the normalized Laplace operators on the whole hypergraph, while the Link bound takes the minimum over the smallest eigenvalues of the adjacency operators on the links. For $2\leq i\leq k-2$, there exists a function f_i supported on the vertices of the link of an $(i-1)$-cell of the complex X and of norm 1 which is an eigenfunction of the normalized adjacency operator on the link with eigenvalue λ_i. Given such f_i, one can define an i-cochain $\widehat{f_i}$ on X supported in the star of an $(i-1)$-cell, such that $\langle\Delta\widehat{f_i},\widehat{f_i}\rangle = 1-\lambda_i$. Since μ_i is the largest eigenvalue of the Laplacian and $\widehat{f_i}$ is of norm 1, $\mu_i\geq 1-\lambda_i$.

4.2. **Sharpness of Hoffman bound on X implies sharpness in $X^{\otimes n}$.** The goal of this section is to prove that the bound 4.1 remains sharp for tensor powers of a hypergraph if it is sharp for the hypergraph itself given all the minimal eigenvalues are negative. First recall the following definition.

Definition 4.2. The tensor product $X\otimes X'$ of two k-uniform hypergraphs $X=(V,\mu)$ and $X'=(V',\mu')$ is a k-uniform hypergraph $(V\times V',\mu\times\mu')$, where $\mu\times\mu'$

stands for the product measure on $(V \times V')^{[k]} \simeq V^{[k]} \times V'^{[k]}$. For a k-uniform hypergraph X, we denote by $X^{\otimes n} = \underbrace{X \otimes \cdots \otimes X}_{n}$ its n-th tensor power.

Proposition 4.3. *Let $X = X \otimes X'$ be the tensor product of k-uniform hypergraphs X and X'. Then for all $0 \leq i < k-2$, the following holds for the smallest eigenvalues of the normalized adjacency operator on the links of its i-faces:*

$$\lambda_i(X) = \min_{\sigma \in X^{(i-1)}} \left[\lambda\left(S\left(X_\sigma\right)\right)\right] = \begin{cases} \lambda_i\left(X\right)\lambda_i\left(X'\right), & \textit{if } \lambda_i\left(X_j\right) \geq 0 \textit{ for } j = 1,2; \\ \min\left\{\lambda_i\left(X\right), \lambda_i\left(X'\right)\right\}, & \textit{otherwise.} \end{cases}$$

In particular, for the tensor power it reads as

$$\lambda_i(X^{\otimes n}) = \begin{cases} \lambda_i\left(X\right), & \textit{if } \lambda_i\left(X\right) \leq 0; \\ \lambda_i\left(X\right)^n, & \textit{if } \lambda_i\left(X\right) \geq 0. \end{cases}$$

Proof. The proposition is a combination of the following two facts. The first is that for an i-face $\sigma = (\sigma_1, \sigma_2) \in X^{(i)}$, its link is the tensor product of the links, i.e.,$X_\sigma = X_{\sigma_1} \otimes X_{\sigma_2}$. The second one is that the matrix of the normalized adjacency operator of the tensor product of two graphs is the Kronecker product of the corresponding matrices of the factors. The result, which dates back to Kronecker himself, states that the eigenvalues of the Kronecker product are exactly the products of the eigenvalues of the factors (see [35] for the proof in the uniform case). Finally, since T_X is a Markov matrix, all of its eigenvalues are bounded in magnitude by 1, which implies the stated formula. □

Theorem 4.4. *Let $X = (V, \mu)$ be a k-uniform hypergraph such that $\lambda_i \leq 0$ for all $0 \leq i < k-2$ and such that the bound (4.1) is sharp for it, i.e.,*

$$\alpha(X) = 1 - \frac{1}{\left(1-\lambda_0\right)\left(1-\lambda_1\right)\cdots\left(1-\lambda_{k-2}\right)},$$

Then it is also sharp for $X^{\otimes n}$ for any positive integer n, and

$$\alpha(X^{\otimes n}) = 1 - \frac{1}{\left(1-\lambda_0\right)\left(1-\lambda_1\right)\cdots\left(1-\lambda_{k-2}\right)}.$$

Proof. It is a direct corollary of Proposition 4.3 that the r.h.s. of the bound (4.1) is the same for $X^{\otimes n}$ as for X. In order to show that it is sharp, note that if $I \subseteq V$ is a maximal independent set in X, that is $\mu(I) = \alpha(X)$, then the set $(I, V, \ldots, V) \subseteq V^n$ is independent in $X^{\otimes n}$. □

4.3. **Computing the generalized Hoffman bound.** Given a k-uniform hypergraph X on the vertex set V and a distribution ν on V, the generalized Hoffman bound gives an upper bound on the ν-measure of an independent set of X for each k-uniform weighted hypergraph $X = (V, \mu)$ whose weight function satisfies the following two constraints: $\mu_1 = \nu$ and $\mu(x) = 0$ whenever $x \notin X$. We can formulate the best bound obtainable in this way as a problem whose variables are the entries of μ:

$$\begin{aligned} &\min\left(1-\lambda_0\right)\cdots\left(1-\lambda_{k-2}\right) \\ &s.t.\, T_{X_s} \succeq \lambda_{|s|}\mathrm{Id} \\ &\quad \mu_1 = \nu \\ &\quad \mu(x) = 0\ \forall x \notin X \\ &\quad \mu \geq 0 \end{aligned}$$

In this program, s goes over all possible faces, $T_{X_s} \succeq \lambda_{|s|}\mathrm{Id}$ means that $T_{X_s} - \lambda_{|s|}\mathrm{Id}$ is positive semidefinite, and $\mu \geq 0$ means that all entries of $\mu$are nonnegative. If a solution to the program has objective value β, then this gives a bound of $1 - 1/\beta$ on the α-measure of an independent set of X.

Since the maximal eigenvalue of T_{X_s} is always 1, we can rephrase the constraint $T_{X_s} \succeq \lambda_{|s|}\mathrm{Id}$ equivalently as follows: the spectral radius of $\mathrm{Id} - T_{X_s}$ is at most $1 - \lambda_{|s|}$. Using Schur complements, this is easily seen to be equivalent to the semidefinite constraint

$$\begin{pmatrix} (1-\lambda_{|s|})\mathrm{Id} & \mathrm{Id} - T_{X_s} \\ \mathrm{Id} - T_{X_s}^T & (1-\lambda_{|s|})\mathrm{Id} \end{pmatrix} \succeq 0.$$

Making $1 - \lambda_{|s|}$ a variable, we have expressed the problem of finding the best generalized Hoffman bound as minimizing a semidefinite program whose objective value is a product of $k-1$ variables. When $k = 2$, this is just a semidefinite program, which can be solved efficiently; up to the nonnegativity constraint $\mu \geq 0$, we have recovered the Lovász θ function. When $k > 2$, the objective function is no longer convex, and it is not clear how to solve the program efficiently.

5. Frankl's problem on extended triangles

5.1. **The uniform version.** Frankl's Turán problem on hypergraphs without extended triangles reads as follows. A *triangle* in $\mathcal{P}([n])$, the power set on $[n]$, is a $2k$-uniform hypergraph supported on three sets $\{A, B, C\}$ such that each element of $[n]$ belongs to an even number of the sets A, B, C. In other words, there exist disjoint k-element sets D, E, F such that $D \cup E = A$, $D \cup F = B$, and $E \cup F = C$. Frankl, [14], asked how large can a family $\mathcal{F} \subseteq \binom{[n]}{2k}$ be if it does not contain a triangle. The reason for considering only even uniformities is that no k-uniform triangle exists for an odd k. Equivalently, we are interested in the maximum independent set in the 3-uniform hypergraph X whose vertices are the $2k$-subsets of $[n]$ and whose hyperedges are triangles.

The skeleton of X is the graph on the same set of vertices whose edges are pairs of subsets whose intersection has size exactly k. This graph is also known as the generalized Johnson graph $J(n, 2k, k)$. Brouwer et al., [5], showed that when $n \leq 4k - 1$, the minimum eigenvalue is

$$\lambda_0 = \frac{n-4k}{2(n-2k)} = 1 - \frac{n}{2(n-2k)}.$$

Since the 3-faces in X are the triples of the form $A, B, A \triangle B$, the links of a vertex is a perfect matching, hence $\lambda_1 = -1$. It follows that when $n \leq 4k - 1$, the size of an independent set is at most

$$\binom{n}{2k}\left(1 - \frac{1}{2(1-\lambda_0)}\right) = \binom{n-1}{2k-1}.$$

In particular, when $n = 4k - 1$, an independent set contains at most a $\frac{2k}{4k-1} = \frac{1}{2} + O(\frac{1}{k})$ fraction of subsets. Now suppose that $n \geq 4k$, and let $\mathcal{F}$ be a triangle-free family. Consider the following random experiment: choose a random $(4k-1)$-subset S of $[n]$, and check whether a random $2k$-subset of S belongs to $\mathcal{F}$. On the one hand, this is at most $\frac{2k}{4k-1}$. On the other hand, this is equal to the density of $\mathcal{F}$. This completes the proof of the following theorem.

Theorem 5.1. *If $\mathcal{F}$ is a family of $2k$-subsets of $[n]$ which does not contain three distinct subsets whose symmetric difference is empty, then $|\mathcal{F}| \leq \binom{n-1}{2k-1}$ if $n \leq 4k-1$, and $|\mathcal{F}| \leq (1/2+O(1/k))\binom{n}{2k}$ otherwise.*

This bound is sharp: If $n \leq 4k-1$ then the family of $2k$-sets containing a fixed element satisfies the condition and has size $\binom{n-1}{2k-1}$. Otherwise, the family of $2k$-sets containing an odd number of elements among the first $\lfloor n/2 \rfloor$ satisfies condition and asymptotically contains half the $2k$-sets.

Frankl, [14], gave the upper bound $(1/2+O(1/n))\binom{n}{k}$, which is slightly better for large n.

5.2. **The p-biased version.** The p-biased version of the problem is as follows:

> Given $p \leq \frac{2}{3}$, how large can $\mu_p(\mathcal{F})$ be if $\mathcal{F} \subseteq \mathcal{P}([n])$ does not contain a triangle?

The reason for the condition $p \leq \frac{2}{3}$ is the fact that the example $\left\{A : |A| > \frac{2}{3}n\right\}$ is triangle-free, and its p-biased measure tends to 1 as n tends to infinity.

The p-biased version of Frankl's problem is the problem of determining the independence number of the 3-uniform hypergraph $X^{\otimes n}$, where $X = (V, \mu)$ is with $V = \{0,1\}$ and

$$\mu([1,1,0]) = \frac{3}{2}p,\ \mu([0,0,0]) = 1 - \frac{3}{2}p.$$

The induced measures are

$$\mu_2([1,1]) = \frac{1}{2}p,\ \mu_2([1,0]) = p,\ \mu_2([0,0]) = 1 - \frac{3}{2}p;$$
$$\mu_1(0) = 1-p,\ \mu_1(1) = p.$$

And therefore, the matrix of the normalized adjacency operator T_X on the skeleton of X is

$$T_X = \begin{pmatrix} \frac{1-\frac{3}{2}p}{1-p} & \frac{\frac{1}{2}p}{1-p} \\ \frac{\frac{1}{2}p}{p} & \frac{\frac{1}{2}p}{p} \end{pmatrix} = \begin{pmatrix} \frac{2-3p}{2(1-p)} & \frac{p}{2(1-p)} \\ \frac{1}{2} & \frac{1}{2} \end{pmatrix},$$

with eigenvalues 1 and $\frac{1-2p}{2(1-p)}$. The induced distribution on the link of $[0]$ is supported on the faces $[0,0]$ and $[1,1]$, hence the corresponding matrix T_{X_0} is

$$T_{X_0} = \begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix},$$

while the link of $[1]$ is supported on $[0,1]$, and T_{X_1} is

$$T_{X_1} = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}.$$

The above shows that $\lambda_0 = \frac{1-2p}{2(1-p)}$ and $\lambda_1 = -1$. When $p > 1/2$, λ_0 is negative, implying

$$\alpha(X^{\otimes n}) \leq 1 - \frac{1}{(1 - \frac{1-2p}{2(1-p)}) \cdot 2} = p.$$

When $p \leq 1/2$, $\lambda_0(X)$ is nonnegative, and so $\lambda_0(X^{\otimes n}) \geq 0$, implying

$$\alpha(X^{\otimes n}) \leq 1 - \frac{1}{1 \cdot 2} = \frac{1}{2}.$$

This completes the proof of the following statement.

Theorem 5.2. *Let $\{0,1\}^n$ denote the space of $\{0,1\}$-vectors of length n, and μ be the p-biased measure on it, with $p \leq 2/3$. If $\mathcal{F} \subseteq \{0,1\}^n$ is a family of vectors which does not contain three distinct vectors whose sum to zero, then $\mu(\mathcal{F}) \leq \max(p, 1/2)$.*

This bound is sharp: if $p \leq 1/2$ then the family of vectors having odd parity satisfies the condition and has measure tending to $1/2$ as $n \to \infty$, and if $p \geq 1/2$ then the set of all vectors having 1 as their first coordinate satisfies the condition and has measure p.

6. Mantel's Theorem

The classical Mantel's theorem bounds the number of edges in a triangle-free graph.

Theorem 6.1. *[29] If a graph on n vertices contains no triangles, then it contains at most $\left\lfloor \frac{n^2}{4} \right\rfloor$ edges.*

We give a spectral proof of Mantel's Theorem for graphs with 2^n vertices that relies on a variation of the bound (4.1). Apart from presenting a spectral proof of this, we would like to show the flexibility of the presented spectral approach. In some cases, one can improve the bound (4.1) by taking not the smallest eigenvalue of the normalized adjacency operator but a larger one. This is possible when the characteristic function of the independent set we are interested in is orthogonal to the eigenfunctions that correspond to smaller eigenvalues.

Proof. First, we encode the statement of the theorem in terms of the independent sets of a hypergraph. Let G be a triangle-free graph on 2^n vertices, which we identify with the set $[2^n]$. Let $\mathcal{X}_{2^n}$ be a 3-partite 3-uniform hypergraph on the vertex set $V = V_1 \cup V_2 \cup V_3$, where each part V_i, $i = 1,2,3$, is a copy of the set $[2^n] \times [2^n]$. The 3-faces of $\mathcal{X}_{2^n}$ are the triples of the form $[(i,j), (j,k), (k,i)]$, where $1 \leq i,j,k \leq 2^n$. Assume the probability distribution on the 3-faces to be the uniform probability distribution on these triples. We encode G as an independent set I of $\mathcal{X}_{2^n}$, namely, $I = I_1 \cup I_2 \cup I_3$, where for each $i = 1,2,3$,

$$I_i = \{(v,u) \in V_i : \text{the set } \{v,u\} \text{ makes an edge of } G\}.$$

Note that $\mathcal{X}_2^{\otimes n} = \mathcal{X}_{2^n}$ and hence this gives the desired encoding of G as the independent set I in a 3-uniform hypergraph which is also a tensor power.

It follows immediately from the construction of $\mathcal{X}_{2^n}$, in particular, from the fact that it is 3-partite, that the matrix of the normalized adjacency operator T on the skeleton of $\mathcal{X}_{2^n}$ is the Kronecker product of the following 3×3 matrix

$$M = \begin{pmatrix} 0 & \frac{1}{2} & \frac{1}{2} \\ \frac{1}{2} & 0 & \frac{1}{2} \\ \frac{1}{2} & \frac{1}{2} & 0 \end{pmatrix}$$

and the matrix M_n which, in turn, is the n-th tensor power of M_1, given by the following 4×4 matrix indexed by the elements of $[2]\times[2]$:

$$M_1 = \begin{array}{c} (1,1)\\ (1,2)\\ (2,1)\\ (2,2) \end{array} \left\{ \underbrace{\begin{pmatrix} \frac{1}{2} & \frac{1}{4} & \frac{1}{4} & 0\\ \frac{1}{4} & 0 & \frac{1}{2} & \frac{1}{4}\\ \frac{1}{4} & \frac{1}{2} & 0 & \frac{1}{4}\\ 0 & \frac{1}{4} & \frac{1}{4} & \frac{1}{2} \end{pmatrix}}_{(1,1)\quad (1,2)\quad (2,1)\quad (2,2)} \right. .$$

Since T is the Kronecker product of M and M_n, its eigenvalues are exactly the products of the eigenvalues of M and M_n. The eigenvectors and eigenvalues of M are

$$\left(\begin{pmatrix}1\\1\\1\end{pmatrix},1\right),\left(\begin{pmatrix}1\\0\\-1\end{pmatrix},-\frac{1}{2}\right),\left(\begin{pmatrix}1\\-1\\0\end{pmatrix},-\frac{1}{2}\right).$$

The eigenvectors and eigenvalues of M_1 are

$$(\chi_1,\lambda_1)=\left(\begin{pmatrix}1\\1\\1\\1\end{pmatrix},1\right),(\chi_2,\lambda_2)=\left(\begin{pmatrix}-1\\1\\1\\-1\end{pmatrix},0\right),$$
$$(\chi_3,\lambda_3)=\left(\begin{pmatrix}1\\0\\0\\-1\end{pmatrix},\frac{1}{2}\right),(\chi_4,\lambda_4)=\left(\begin{pmatrix}0\\1\\-1\\0\end{pmatrix},-\frac{1}{2}\right).$$

We now exploit the symmetries of the set I to show that its characteristic function 1_I is orthogonal to the subspace of eigenvectors of T with negative eigenvalues. First, note that the set I is invariant under the action of the symmetric group S_3 acting on $\mathcal{X}_{2^n}$ by permutations of the parts $\{V_1,V_2,V_3\}$. The only eigenvector of M invariant under the action of S_3 is the constant vector with eigenvalue 1. A vertex in $\mathcal{X}_2$ is a pair (i,j). Let $\mathcal{S}_0$ be the operator that swaps between i and j, which satisfies

$$\mathcal{S}_0\chi_i=\begin{cases}\chi_i & i=1,2,3\\ -\chi_i & i=4\end{cases}.$$

The operator $\mathcal{S}$ that swaps the coordinates of a vertex in $\mathcal{X}_{2^n}$ is of the form $\mathcal{S}=\mathcal{S}_0^{\otimes n}$. The set I is invariant under the action of $\mathcal{S}$ by the construction.

Since the characteristic function 1_I is orthogonal to the subspace spanned by eigenvectors of the normalized adjacency operator with negative eigenvalues, we may take 0 instead of λ_0 in (4.1). Note that since the link of most vertices in $\mathcal{X}_{2^n}$ is bipartite, $\lambda_1=-1$ (which is the minimal possible eigenvalue of a Markov matrix). Hence, the bound (4.1) reads as

$$\frac{|I|}{3\cdot 4^n}\le 1-\frac{1}{1\cdot 2}=\frac{1}{2}.$$

Taking into account the fact that every edge of G is counted six times in $|I|$ completes the proof. $\square$

7. Frankl-Tokushige Theorem on Intersecting Families

Our method also provides a new proof for the result of Frankl and Tokushige on k-wise intersecting families, [16].

Theorem 7.1. *[16] Let $k \geq 2$ and $p \leq 1 - \frac{1}{k}$. Assume $\mathcal{F} \subset \mathcal{P}([n])$ is k-wise intersecting, that is, for all $F_1, \ldots, F_k \in \mathcal{F}$*

$$F_1 \cap \cdots \cap F_k \neq \emptyset.$$

Then $\mu_p(\mathcal{F}) \leq p$, where μ_p stands the p-biased measure.

Proof. Let X be the k-uniform hypergraph on $\{0,1\}$ weighted by the measure $\mu([0^{(k)}]) = 1 - \frac{k}{k-1}p$, $\mu([0, 1^{(k-1)}]) = \frac{k}{k-1}p$. Here $0^{(k)}$ means k copies of 0. The induced measure μ_1 on the vertex set $\{0,1\}$ is the p-biased one, i.e., $\mu_1(1) = p$ and $\mu_1(0) = 1 - p$. The matrix form of T_X is directly calculated to be

$$\begin{pmatrix} \frac{1-\frac{k}{k-1}p}{1-p} & \frac{\frac{1}{k-1}p}{1-p} \\ \frac{1}{k-1} & \frac{k-2}{k-1} \end{pmatrix},$$

from which it follows that its smallest eigenvalue is $\lambda_0 = \frac{\frac{k-2}{k-1}-p}{1-p} < 0$ hence $\frac{1}{1-\lambda_0} = (k-1)(1-p)$.

In order to calculate λ_i for $i > 0$, notice that the only links with non-trivial $\mu_1(X_S)$ are of the faces $S = [1^\ell]$. The matrix form of T_{X_S} is directly calculated to be

$$\begin{pmatrix} 0 & 1 \\ \frac{1}{k-1-\ell} & \frac{k-2-\ell}{k-1-\ell} \end{pmatrix},$$

and so $\lambda_\ell = -\frac{1}{k-1-\ell} < 0$ and $\frac{1}{1-\lambda_\ell} = \frac{k-1-\ell}{k-\ell}$. Applying the generalized Hoffman bound for tensor powers, we conclude

$$\alpha(X^{\otimes n}) \leq 1 - (k-1)(1-p) \cdot \frac{k-2}{k-1} \cdot \cdots \cdot \frac{1}{2} = p.$$

Since the edges of X are exactly the multisets with either all 0's or exactly one 0, every k-wise intersecting set is independent in $X^{\otimes n}$. □

Acknowledgements

The authors are grateful to Ehud Friedgut, Gil Kalai, Guy Kindler, and Dor Minzer for valuable discussions.

Funding. YF is a Taub Fellow, and is supported by the Taube Foundation and ISF grant 1337/16. KG is currently supported by the SNF grant number 20002_169106, and by ERC grant 336283 while at the Weizmann Institute and Bar-Ilan University.

References

[1] Rudolf Ahlswede and Gyula OH Katona. Contributions to the geometry of hamming spaces. *Discrete Mathematics*, 17(1), 1977.

[2] Noga Alon, Irit Dinur, Ehud Friedgut, and Benny Sudakov. Graph products, fourier analysis and spectral techniques. *Geometric & Functional Analysis GAFA*, 14(5):913–940, 2004.

[3] Noga Alon and Eyal Lubetzky. Independent sets in tensor graph powers. *Journal of Graph Theory*, 54(1):73–87, 2007.

[4] Christine Bachoc, Anna Gundert, and Alberto Passuello. The theta number of simplicial complexes. *Israel Journal of Mathematics*, 232(1):443–481, 2019.

[5] Andries E Brouwer, Sebastian M Cioabă, Ferdinand Ihringer, and Matt McGinnis. The smallest eigenvalues of hamming graphs, johnson graphs and other distance-regular graphs with classical parameters. *Journal of Combinatorial Theory, Series B*, 133:88–121, 2018.

[6] Irit Dinur and Ehud Friedgut. Intersecting families are essentially contained in juntas. *Combinatorics, Probability and Computing*, 18(1-2):107–122, 2009.

[7] Irit Dinur, Ehud Friedgut, and Oded Regev. Independent sets in graph powers are almost contained in juntas. *Geometric and Functional Analysis*, 18(1):77–97, 2008.

[8] Irit Dinur, Elchanan Mossel, and Oded Regev. Conditional hardness for approximate coloring. *SIAM Journal on Computing*, 39(3):843–873, 2009.

[9] Irit Dinur and Samuel Safra. On the hardness of approximating minimum vertex cover. *Annals of mathematics*, pages 439–485, 2005.

[10] David Ellis, Yuval Filmus, and Ehud Friedgut. Triangle-intersecting families of graphs. *Journal of the European Mathematical Society*, 14:841–885, 2012.

[11] David Ellis, Ehud Friedgut, and Haran Pilpel. Intersecting families of permutations. *Journal of the American Mathematical Society*, 24(3):649–682, 2011.

[12] Paul Erdős. A problem on independent r-tuples. *Ann. Univ. Sci. Budapest*, 8:93–95, 1965.

[13] Paul Erdős, Chao Ko, and Richard Rado. Intersection theorems for systems of finite sets. *The Quarterly Journal of Mathematics*, 12(1):313–320, 1961.

[14] P Frankl. Asymptotic solution of a Turán-type problem. *Graphs and Combinatorics*, 6(3):223–227, 1990.

[15] Péter Frankl. On Sperner families satisfying an additional condition. *Journal of Combinatorial Theory, Series A*, 20(1):1–11, 1976.

[16] Péter Frankl and Norihide Tokushige. Weighted multiply intersecting families. *Studia Scientiarum Mathematicarum Hungarica*, 40(3):287–291, 2003.

[17] Ehud Friedgut. Boolean functions with low average sensitivity depend on few coordinates. *Combinatorica*, 18(1):27–35, 1998.

[18] Ehud Friedgut. On the measure of intersecting families, uniqueness and stability. *Combinatorica*, 28(5):503–528, 2008.

[19] Ehud Friedgut and Oded Regev. Kneser graphs are like swiss cheese. *Discrete Analysis*, (2):18pp., 2018.

[20] Konstantin Golubev. On the chromatic number of a simplicial complex. *Combinatorica*, pages 1–12, 2016.

[21] Konstantin Golubev and Ori Parzanchevski. Spectrum and combinatorics of two-dimensional ramanujan complexes. *Israel Journal of Mathematics*, 230(2):583–612, 2019.

[22] Anna Gundert and May Szedlák. Higher dimensional discrete cheeger inequalities. *JoCG*, 6:54–71, 2015.

[23] Alan J Hoffman. On eigenvalues and colorings of graphs. In *Selected Papers Of Alan J Hoffman: With Commentary*, pages 407–419. World Scientific, 2003.

[24] Danijela Horak and Jürgen Jost. Spectra of combinatorial laplace operators on simplicial complexes. *Advances in Mathematics*, 244:303–336, 2013.

[25] Noam Lifshitz. Nearly independent sets in tensor powers of simplicial complexes. in preparation.

[26] Nathan Linial and Yuval Peled. On the phase transition in random simplicial complexes. *Annals of Mathematics*, pages 745–773, 2016.

[27] László Lovász. On the shannon capacity of a graph. *IEEE Transactions on Information theory*, 25(1):1–7, 1979.

[28] Alexander Lubotzky. Ramanujan complexes and high dimensional expanders. *Japanese Journal of Mathematics*, 9(2):137–169, 2014.

[29] W. Mantel. Problem 28 (solution by h. gouwentak, w. mantel, j. teixeira de mattes, f. schuh and w. a. wythoff). *Wiskundige Opgaven*, 10:60–61, 1907.

[30] Roy Meshulam. On subsets of finite abelian groups with no 3-term arithmetic progressions. *Journal of Combinatorial Theory, Series A*, 71(1):168–172, 1995.

[31] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz. Noise stability of functions with low influences: Invariance and optimality. *Annals of Mathematics*, pages 295–341, 2010.

[32] Ori Parzanchevski. Mixing in high-dimensional expanders. *Combinatorics, Probability and Computing*, 26(5):746–761, 2017.

[33] Ori Parzanchevski and Ron Rosenthal. Simplicial complexes: spectrum, homology and random walks. *Random Structures & Algorithms*, 50(2):225–261, 2017.

[34] Ori Parzanchevski, Ron Rosenthal, and Ran J Tessler. Isoperimetric inequalities in simplicial complexes. *Combinatorica*, 36(2):195–227, 2016.
[35] Paul M Weichsel. The kronecker product of graphs. *Proceedings of the American mathematical society*, 13(1):47–52, 1962.
[36] Richard M Wilson. The exact bound in the Erdös-Ko-Rado theorem. *Combinatorica*, 4(2-3):247–257, 1984.

(Y. Filmus) yuvalfi@cs.technion.ac.il, Technion – Israel Institute of Technology, Haifa, Israel.

(K. Golubev) golubevk@ethz.ch, D-MATH, ETH Zurich, Switzerland.

(N. Lifshitz) noam.lifshitz@gmail.com, Einstein Institute of Mathematics, Hebrew University, Jerusalem, Israel.