ALGEBRAIC COMBINATORICS

Yuval Filmus, Konstantin Golubev & Noam Lifshitz

High dimensional Hoffman bound and applications in extremal combinatorics

High dimensional Hoffman bound and applications in extremal combinatorics

Yuval Filmus, Konstantin Golubev & 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. 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. In this paper we introduce the problem of upper-bounding independent sets in tensor powers of hypergraphs. We show that many prominent open problems in extremal combinatorics, such as the Turán problem for graphs and hypergraphs, can be encoded as special cases of this problem. We generalize the Hoffman bound to hypergraphs, and give several applications.

1. INTRODUCTION

The celebrated Hoffman bound [30] 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.[1] The Hoffman bound, in a generalized version due to Lovász [33], 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, as well as many other Erdős–Ko–Rado theorems on various domains [12, 13, 25]. 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 [17], 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.

[1]Hoffman's original paper only states a bound on the chromatic number. The corresponding bound on the independence number appears in works of Haemers [29] and Lovász [33], where it is attributed to Hoffman. The weighted version of the Hoffman bound, used by Lovász [33] to compute the Shannon capacity of the pentagon, had appeared earlier in foundational work of Delsarte [7] in coding theory.

YUVAL FILMUS, KONSTANTIN GOLUBEV & NOAM LIFSHITZ

Schrijver [44] generalized the Hoffman bound to handle certain conditions on triples of sets, in the context of coding theory. His generalization can handle constraints on the possible values of $(|A\Delta B|, |A\Delta C|, |B\Delta C|)$, where A, B, C are sets in the family. His method has been used in several subsequent works, for example [6, 23, 24].

Recently, the Hoffman bound has been generalized to hypergraphs [4, 26]. Our 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 $[v_1, \ldots, v_i]$. The collection $V^{[0]}$ consists of 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 $0 \leqslant i \leqslant k-1$, define a probability measure μ_i (or $\mu_i(X)$, if we wish to make the hypergraph explicit) 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 $X^{(0)} = \{\varnothing\}$, i.e. the empty set is the one and only 0-face of X. We assume without loss of generality that $X^{(1)} = V$, that is, $\mu_1(v) > 0$ for all $v \in V$ (otherwise, we can replace V with $X^{(1)}$).

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, 26]) 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.

We caution the reader that in the literature on abstract simplicial complexes, an i-face is usually defined as a face of size $i+1$.

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 $\mu_1(I)$, 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 $\mu_1(I) = \alpha(X)$.*

A *flag* of X is a tuple of faces $A_0, \ldots, A_k$ such that A_i is an i-face and $A_0 \subseteq \cdots \subseteq A_k$. 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 an order of the vertices it contains uniformly at random.

We use the notation $\mathbf{v} \sim \mu_1$ to denote a random variable $\mathbf{v}$ supported on $V^{[1]}$ and distributed according to μ_1. Similarly, $(\mathbf{u}, \mathbf{v}) \sim \tilde{\mu}_2$ denotes a random variable $(\mathbf{u}, \mathbf{v})$ supported on $V^{[2]}$ and distributed according to $\tilde{\mu}_2$, and so on.

The notation $\mathbf{w} \sim \boldsymbol{\mu}$ denotes a random variable $\mathbf{w}$ supported on X^k and corresponding to our earlier definition (a probability distribution over tuples of faces $A_0 \subseteq \cdots \subseteq A_k$) as follows: $A_i = [\mathbf{w}_1, \ldots, \mathbf{w}_i]$.

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 $\boldsymbol{\mu}_k \smallsetminus \sigma$. Note that the link of the empty set is the whole hypergraph X itself. For an element $v \in X$, the link of v is $X_v = X_{\{v\}}$.

For a set $A \in X^{(k-i)}$ we shall say that $\mu_\sigma(A)$ is the *relative measure of A according to σ*.

The *skeleton* (or *underlying graph*) 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

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

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

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

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

$$(1) \qquad T_X(u,v) = \begin{cases} \frac{\mu_2([u,u])}{\mu_1(u)} & \text{if } u = v; \\ \frac{\mu_2([u,v])}{2\mu_1(u)} & \text{if } u \neq v. \end{cases}$$

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 $0 \leqslant i \leqslant k-2$, 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 following measure on $(V \times V')^{[k]} \simeq V^{[k]} \times V'^{[k]}$: an ordering of an edge $(\sigma, \sigma') \in V^{[k]} \times V'^{[k]}$ implies an ordering on σ and on σ', we define $(\mu \times \mu')_k(\sigma, \sigma')$ as the sum of $\widetilde{\mu_k}(\sigma) \widetilde{\mu'_k}(\sigma')$ over all orderings of (σ, σ'). Recall that $\widetilde{\mu_k}$ is constructed from μ by choosing an ordering on an edge uniformly at random. 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*

$$(2) \qquad \alpha(X) \leqslant 1 - \frac{1}{(1 - \lambda_0)(1 - \lambda_1) \cdots (1 - \lambda_{k-2})}.$$

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

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

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

Furthermore, if $\lambda_0 > -1$ and I is an independent set attaining the bound, then I is a dictator (viewed as a subset of V^n, membership in I depends on a single coordinate).

We prove the first part of this theorem as Theorem 4.2, and the second part (regarding $X^{\otimes n}$) as Theorem 4.6. The proof also yields conditions for equality, which we use in the applications to identify extremal families. Since the condition for equality in the first part of the theorem is somewhat verbose, we only state it in full after proving Theorem 4.2.

When $k = 2$, we obtain the classical Hoffman bound $\alpha(X) \leqslant \frac{-\lambda}{1-\lambda}$, where λ is the minimal eigenvalue of a weighted adjacency matrix of X.

We apply Theorem 1.4 to deduce the following result on 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 $3k \leqslant n \leqslant 4k-1$. 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 $|\mathcal{F}| \leqslant \binom{n-1}{2k-1}$. 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 $1/2 \leqslant p \leqslant 2/3$. If $\mathcal{F} \subseteq \{0,1\}^n$ is a family of vectors which does not contain three distinct vectors whose sum is zero, then $\mu(\mathcal{F}) \leqslant p$. 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 proofs of Mantel's theorem on triangle-free graphs and Frankl–Tokushige theorem on k-wise intersecting families.

THEOREM 1.6 (Mantel [35]). *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 (Frankl–Tokushige [19]). *Let $k \geqslant 2$ and $p \leqslant 1 - \frac{1}{k}$. Assume $\mathcal{F} \subset \mathcal{P}([n])$ is k-wise intersecting family of subsets of $[n]$, that is, for all $F_1, \ldots, F_k \in \mathcal{F}$*

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

Then $\mu_p(\mathcal{F}) = \sum_{F \in \mathcal{F}} p^{|F|}(1-p)^{n-|F|} \leqslant p$, where 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) 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.[2]

THEOREM 2.1 (Hoffman bound [29, 30, 33]). *Let G be a graph. Then*
$$\alpha(G) \leqslant \frac{-\lambda_{\min}(G)}{1 - \lambda_{\min}(G)}.$$

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 [33] showed that it can be used to prove the celebrated Erdős–Ko–Rado theorem [15], (aka EKR theorem), which states that for $n \geqslant 2k$, 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, Wilson [47] 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 the Hoffman 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, $n \geqslant (t+1)(k-t+1)$.

This approach became more systematic in the work of Friedgut [21], 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 [12]. Ellis, Friedgut, and Pilpel [13], 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). For an exposition of many different applications of the Hoffman bound to Erdős–Ko–Rado theory, consult the excellent monograph of Godsil and Meagher [25].

Another important application of this method is to error-correcting codes. This approach was pioneered by Delsarte [7]. Schrijver [43] showed that Delsarte's linear programming bound can be expressed as a mild generalization of the Hoffman bound.

---

[2] The Hoffman bound is sometimes stated for d-regular graphs in terms of the unnormalized independence number and the unnormalized adjacency matrix, in which case it reads $\alpha(G) \leqslant \frac{-\lambda_{\min}(G)}{d - \lambda_{\min}(G)} n$, where n is the number of vertices. This formula can be derived from ours by noticing that normalizing the adjacency matrix of a d-regular graph corresponds to division by d.

McEliece, Rodemich, Rumsey, and Welch [36] used this approach to prove the so-called "linear programming bound", which is the best bound on the size of binary error-correcting codes. (Since then, the bound has been recovered using different techniques by Navon and Samorodnitsky [39].)

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.

Schrijver [44] came up with a generalization of the Hoffman bound which is able to handle some constraints on triples of sets. In more detail, his method can handle constraints on the allowed Hamming distance patterns on triples of sets in the family. However, his method does not handle arbitrary constraints on triples, and does not extend beyond triples.

Over the years there has been great interest in the problem of generalizing results from graphs to hypergraphs (or, equivalently, simplicial complexes), see e.g. [27, 28, 32, 34, 40, 41, 42]. In particular, two generalizations of the Hoffman bound were given by [4] and [26]. Such results are known as *high-dimensional Hoffman bounds*. In contrast to Schrijver's generalization, such bounds apply to arbitrary hypergraphs. 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 [9]. One motivation is their connection to the theory of hardness of approximation, see [10]. However, our main focus in this paper will be the connection to extremal set theory. This connection was first implicitly established by Friedgut [21], and later more explicitly by Dinur and Friedgut [8].

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 $p = \frac{2m}{n(n-1)}$.[3] 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?

---

[3]Both models construct random graphs on n vertices. In the model $G(n, p)$, we include each edge independently with probability p. In the model $G(n, m)$, we choose exactly m edges at random.

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 Dinur and Safra [11].

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 [20], Dinur and Friedgut [11], and later by Friedgut and Regev [22], 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: $\lambda_{\min}(G^{\otimes n}) = \lambda_{\min}(G)$. This immediately implies that the Hoffman bound is sharp on $G^{\otimes n}$ whenever it is sharp on G. It reduces the p-biased EKR problem for families $\mathcal{F} \subseteq \mathcal{P}([n])$ 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 $1 \leqslant i \leqslant n$ 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], essentially 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. $\alpha(G) = \frac{-\lambda_{\min}}{1 - \lambda_{\min}}$, then $\alpha(G^{\otimes n}) = \alpha(G)$. Moreover, if A is an independent set with $\mu_G(A) = \alpha(G)$, then A is a dictatorship.*

Furthermore, for each $\epsilon > 0$ there exists $\delta = \delta(\epsilon, G) > 0$ such that if an independent set A satisfies $\mu_G(A) > \alpha(G) - \delta$, then there exists an independent dictatorship B such that $\mu_G(A \Delta B) < \epsilon$.

Moreover, δ in fact depends only on ϵ and on $\min_v \mu_G(\{v\})$.

(To deduce this theorem from their work requires the FKN theorem for product spaces, due to Jendrej, Oleszkiewicz, and Wojtaszczyk [31].)

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, simplicial complexes). In addition to the definitions given in the introduction, let us give a version of them in the k-partite setting, which is a special case of k-uniform hypergraphs.

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 $X^{\otimes n} = (V_1^n, \ldots, V_k^n, \mu^{\otimes n})$, 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 [37], 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 m-uniform 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 [38].

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 (in preparation by the third author). Let us state it here 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 3-uniform hypergraph

$$X^{\otimes n}$$

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

$$G_1 \subseteq \mathcal{A}_1 \times \mathcal{A}_2, G_2 \subseteq \mathcal{A}_2 \times \mathcal{A}_3, G_3 \subseteq \mathcal{A}_3 \times \mathcal{A}_1,$$

where each $\mathcal{A}_i$ 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 [14], 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 $p \leqslant \frac{1}{s}$, how large can $\mu_p\left(\mathcal{F}\right)$ 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 $\mu\left([0, \ldots, 0]\right) = 1 - sp$ (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 [17]. 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 $k \leqslant \frac{s-1}{s}n$, 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 [19]. They showed that the largest value of $\mu_p\left(\mathcal{F}\right)$ for an s-wise intersecting family $\mathcal{F}$ is p, as long as $p \leqslant \frac{s-1}{s}$. This problem can be expressed as the determining the independence number of the s-uniform 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) $\mu\left(x\right) = 0$ if x is the all ones vector.

It is easy to construct many hypergraphs X that satisfy these hypotheses. We reprove this result in Section 7.

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 $\mathcal{P}\left([n]\right)$ 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 [18] 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 $p \leqslant \frac{2}{3}$, how large can $\mu_p\left(\mathcal{F}\right)$ be if $\mathcal{F} \subseteq \mathcal{P}\left([n]\right)$ does not contain a triangle?

The reason for the condition $p \leqslant \frac{2}{3}$ 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 3-uniform hypergraph $X^{\otimes n}$, where $X = (V, \mu)$ is with $V = \{0, 1\}$ and

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

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, where $n \geqslant 3k$, then the family contains at most $\frac{2k}{\min(n, 4k-1)} \binom{n}{2k}$ subsets.*

Furthermore, when $n \leqslant 4k - 1$ and $p \geqslant 1/2$ 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 $p \leqslant 2/3$, 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 simplicial 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 by the second author in [26], and by Bachoc, Gundert, and Passuello in [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 [26].

Let $X = (V, \mu)$ be a k-uniform hypergraph on the vertex set V. Recall (see Subsection 1.1) that for $0 \leqslant i \leqslant k - 2$ 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,
$$\lambda_i(X) = \min_{\sigma \in X^{(i)}} \left[\lambda(S(X_\sigma)) \right].$$

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 4.1. Let $X = (\{0,1\}, \mu)$ be a graph (in other words, a 2-uniform hypergraph) on two vertices $\{0, 1\}$ with the probability measure μ defined on the edges as
$$\mu(\{0,0\}) = p_1, \ \mu(\{0,1\}) = p_2, \ \mu(\{1,1\}) = p_3.$$

Then the induced distribution μ_1 on the vertices is as follows:
$$\mu_1(0) = p_1 + \frac{1}{2}p_2, \ \mu_1(1) = \frac{1}{2}p_2 + p_3.$$

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

(3) $$\alpha(X) \leqslant 1 - \frac{1}{(1 - \lambda_0)(1 - \lambda_1) \cdots (1 - \lambda_{k-2})}.$$

Proof. The proof goes by induction on the uniformity of the hypergraph. The base case, $k = 2$, is the graph case, and the bound (3) 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 $v_1 = 1, v_2, \ldots, v_m$ be an orthonormal basis of its eigenvectors with eigenvalues $1 \geqslant l_2 \geqslant \cdots \geqslant l_m = \lambda_0$ (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,

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

or in other words, it is equal to the probability of an ordered edge (2-face) distributed according to μ_2 to have both ends in I. On the other hand,

$$\langle T_X f, f \rangle = \sum_{i=1}^{m} l_i \langle f, v_i \rangle^2$$

$$\geqslant \langle f, 1 \rangle^2 + \sum_{i=2}^{m} \lambda_0 \langle f, v_i \rangle^2$$

$$= \langle f, 1 \rangle^2 (1 - \lambda_0) + \lambda_0 \langle f, f \rangle$$

$$= \mathbb{E}[f]^2 (1 - \lambda_0) + \lambda_0 \mathbb{E}[f^2].$$

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

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

Note that

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

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

$$\Pr_{\boldsymbol{y} \sim \mu_1(X_x)} [\boldsymbol{y} \in I] \leqslant 1 - \frac{1}{(1 - \lambda_0(X_x)) \cdots (1 - \lambda_{k-3}(X_x))}.$$

Clearly $\lambda_i(X_x) \geqslant \lambda_{i+1}$. Since the expression on the right-hand side is monotone decreasing in $\lambda_i(X_x)$,

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

Combining the above bounds, we deduce that

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

which proves the required bound after rearrangement. $\qquad\square$

Suppose now that an independent set I attains the bound in Theorem 4.2. Then the following inequality involving $f = 1_I$ is tight:

$$\langle T_X f, f \rangle = \sum_{i=1}^{m} l_i \langle f, v_i \rangle^2 \geqslant \langle f, 1 \rangle^2 + \sum_{i=2}^{m} \lambda_0 \langle f, v_i \rangle^2,$$

and so $f - \mathbb{E}[f]$ belongs to the eigenspace of λ_0. Furthermore, for each $x \in I$, the bound

$$\Pr_{\boldsymbol{y} \sim \mu_1(X_x)}[\boldsymbol{y} \in I] \leqslant 1 - \frac{1}{(1 - \lambda_0(X_x)) \cdots (1 - \lambda_{k-3}(X_x))}$$

is tight, and $\lambda_i(X_x) = \lambda_{i+1}$ for $0 \leqslant i \leqslant k - 3$. We obtain the following tightness condition.

THEOREM 4.3. *If I is an independent set attaining the bound in Theorem 4.2 then for every multiset $A \subseteq I$ of size $\ell \leqslant k - 2$:*

- $\lambda_i(X_A) = \lambda_{i+\ell}$ *for $0 \leqslant i \leqslant k - 2 - \ell$. In particular, $\lambda_0(X_A) = \lambda_\ell$.*
- $f - \mu_A(I)$ *lies in the λ_ℓ-eigenspace of T_{X_A}.*

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 [26], 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 4.2 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. However the comparison between them does not seem to be straightforward. It is not clear how the spectra of the Laplace operators relate to the spectra of the normalized adjacency operators of the links used to prove the Link bound in the general case.

The Theta bound follows a different approach. In [4], the authors show that the Theta bound is incomparable to the Laplacian bound by providing four families of simplicial complexes, on half of them one bound is sharp while the other is not, and vice versa.

The Link bound is sharp for all four examples provided in [4], as we show below. In other words, the Link bound performs at least as good, and sometimes better, as the Theta bound and the Laplacian bound. However, we do not have a proof that it is always the case in general.

Here are the examples considered in [4]:

(1) Let S be a set of $3m$ elements of three colors, each color appearing m times. What is the largest size of a subset of S in which no triple of elements contains all three colors? In other words, what is the largest independent set of the 3-partite 3-uniform hypergraph on $3m$ vertices. The answer is $2m$. This is shown by the Theta bound by not by the Laplacian bound.

   The Link bound is also sharp: if we introduce the uniform measure on the hypergraph, then $\lambda_0 = -\frac{1}{2}$ and $\lambda_1 = -1$, and the Link bound reads as $\frac{2}{3}$, which is exactly the measure of the set of $2m$ vertices.

(2) Consider the set S again. What is the largest size of a subset of S in which each triple of elements contains all three colors? The answer is 3. This is again shown by the Theta bound by not by the Laplacian bound.

   The Link bound is also sharp. To show this, we consider the 3-uniform hypergraph whose triangles are all triples of distinct elements of S in which not all colors are the same. We introduce a measure which gives each monochromatic triangle a probability p_1 and each bichromatic triangle a probability p_2, where $p_1 = 2p_2$. We compute $\lambda_0 = -\frac{1}{2m-1}$ and $\lambda_1 = -\frac{1}{2m-2}$, and so the Link bound reads as $\frac{1}{m} = \frac{3}{3m}$.

(3) Let T be a set of $2m$ elements of two colors, each color appearing m times. What is the largest size of a subset of T in which no triple of elements contains

both colors? The answer is m. This is shown by both the Theta bound and the Laplacian bound, and so also by the Link bound.

   The Link bound is also sharp: if we introduce the uniform measure on the 3-uniform hypergraph whose triangles are all bichromatic triples of distinct elements of T, then $\lambda_0 = -\frac{1}{3}$ and $\lambda_1 = -\frac{1}{2}$, and the Link bound reads as $\frac{1}{2}$, which is exactly the measure of the set of m vertices.

(4) Consider the set T again. What is the largest size of a subset of T in which each triple of elements contains both colors? The answer is 4. This is shown by the Laplacian bound and so also by the Link bound, but not by the Theta bound.

   The Link bound is also sharp: if we introduce the uniform measure on the 3-uniform hypergraph whose triangles are all monochromatic triples of distinct elements of T, then $\lambda_0 = -\frac{1}{m-1}$ and $\lambda_1 = -\frac{1}{m-2}$, and so the Link bound reads as $\frac{2}{m} = \frac{4}{2m}$.

More details on these examples can be found in Appendix A.

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 (3) 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 4.4. *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 4.5. *Let $Y = X \otimes X'$ be the tensor product of k-uniform hypergraphs X and X'. Then for all $0 \leqslant i < k - 2$, the following holds for the smallest eigenvalues of the normalized adjacency operator on the links of its i-faces:*

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

In particular, for the tensor power it reads as

$$\lambda_i(X^{\otimes n}) = \begin{cases} \lambda_i \left(X \right), & \text{if } \lambda_i \left(X \right) \leqslant 0; \\ \lambda_i \left(X \right)^n, & \text{if } \lambda_i \left(X \right) \geqslant 0. \end{cases}$$

Proof. The proposition is a combination of the following two facts. The first is that for an i-face $\sigma = (\sigma_1, \sigma_2) \in Y^{(i)}$, its link is the tensor product of the links, i.e. $Y_\sigma = X_{\sigma_1} \otimes X'_{\sigma_2}$. 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 [46] 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. $\square$

THEOREM 4.6 (Restatement of second part of Theorem 1.4). *Let $X = (V, \mu)$ be a k-uniform hypergraph such that $\lambda_i \leqslant 0$ for all $0 \leqslant i < k-2$ and such that the bound (3) 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

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

Furthermore, if $\lambda_0 > -1$ then an independent set attaining the bound is a dictatorship (depends on at most one of the n input coordinates).

Proof. It is a direct corollary of Proposition 4.5 that the r.h.s. of the bound (3) 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 $X^{\otimes n}$.

Suppose now that I is an independent set attaining the bound. Theorem 4.3 shows that $1_I - \mu(I)$ lies in the λ_0-eigenspace of $T_{X^{\otimes n}}$, which consists of functions depending on a single coordinate. Therefore we can express $f = 1_A$ in the form

$$f(x_1, \ldots, x_n) = \sum_{i=1}^n \phi_i(x_i).$$

Since f is Boolean, at most one ϕ_i can be non-constant, and so f is a dictator. $\square$

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:

$$\min (1 - \lambda_0) \cdots (1 - \lambda_{k-2})$$
$$s.t.\, T_{X_s} \succeq \lambda_{|s|} \, \mathrm{Id}$$
$$\mu_1 = \nu$$
$$\mu(x) = 0 \,\, \forall x \notin X$$
$$\mu \geqslant 0$$

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 $\mu \geqslant 0$ means that all entries of μ 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 $\mu \geqslant 0$, 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 [18], 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, Cioabă, Ihringer and McGinnis [5, Theorem 3.10] showed that when $3k \leqslant n \leqslant 4k - 1$, 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 link of a vertex is a perfect matching, hence $\lambda_1 = -1$. It follows that when $3k \leqslant n \leqslant 4k - 1$, 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 $\frac{2k}{4k-1}$ fraction of subsets.

Furthermore, if I is an independent set of maximal size and $f = 1_I$ then $f - \mathbb{E}[f]$ must belong to the eigenspace of λ_0. When $3k \leqslant n < 4k - 1$, this eigenspace consists of all functions of degree 1 (this follows from the techniques of Brouwer et al.), and so $\deg f \leqslant 1$, implying that f is a dictatorship (see [16]).[4]

Now suppose that $n \geqslant 4k$, 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 $|\mathcal{F}| \leqslant \binom{n-1}{2k-1}$ if $3k \leqslant n \leqslant 4k - 1$, and $|\mathcal{F}| \leqslant (\frac{1}{2} + \frac{1}{8k-2})\binom{n}{2k}$ if $n \geqslant 4k$.*

This bound is sharp: If $3k \leqslant n \leqslant 4k - 1$ then the family of $2k$-sets containing a fixed element satisfies the condition and has size $\binom{n-1}{2k-1}$. Furthermore, these are the unique examples if $n < 4k - 1$. If $n \geqslant 4k$ then 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 [18] gave the upper bound $(\frac{1}{2} + \frac{2k}{n-4k})\binom{n}{2k}$, which is better when $n > (4k)^2$.

5.2. THE p-BIASED VERSION. The p-biased version of the problem is as follows:

> Given $p \leqslant \frac{2}{3}$, 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 $p \leqslant \frac{2}{3}$ is the fact that the example $\{A : |A| > \frac{2}{3}n\}$ is triangle-free, and its p-biased measure tends to 1 as n tends to infinity.

---

[4]When $n = 4k - 1$, the methods of Brouwer et al. only imply that $\deg f \leqslant 2$. Indeed, there are other examples: all sets intersecting $\{i, j\}$ at exactly one point. We thank Ferdinand Ihringer for this example.

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

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

The induced measures are

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

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

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

Moreover, $\lambda_0 > -1$, and so all extremal families are dictators.

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

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

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 $p \leqslant 2/3$. If $\mathcal{F} \subseteq \{0, 1\}^n$ is a family of vectors which does not contain three distinct vectors whose sum to zero, then $\mu(\mathcal{F}) \leqslant \max(p, 1/2)$.*

This bound is sharp: if $p \leqslant 1/2$ then the family of vectors having odd parity satisfies the condition and has measure tending to $1/2$ as $n \to \infty$, and if $p > 1/2$ then the set of all vectors having 1 as their first coordinate satisfies the condition and has measure p, and moreover these are the unique extremal families.

6. Mantel's Theorem

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

THEOREM 6.1 ([35]). *If a graph on n vertices contains no triangles, then it contains at most $\left\lfloor \frac{n^2}{4} \right\rfloor$ edges.*

Mantel's theorem has many different proofs. We give a new proof of Mantel's theorem for graphs with 2^n vertices that relies on a variation of the bound (3). Ours is not the first spectral proof of Mantel's theorem: Tait and Tobin [45] give a very simple such proof. However, their proof is very similar to existing arguments, such as the ones that can be found in Zhao's lecture notes [48]. In contrast, our proof is rather different from all existing ones.

Apart from presenting a "genuine" spectral proof of this fundamental result, we would like to show the flexibility of the presented spectral approach. In some cases, one can improve the bound (3) 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 $1 \leqslant i, j, k \leqslant 2^n$. 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 $[2] \times [2]$:

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

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

$$(\chi_1, \lambda_1) = \left(\begin{pmatrix} 1 \\ 1 \\ 1 \\ 1 \end{pmatrix}, 1 \right), \qquad (\chi_2, \lambda_2) = \left(\begin{pmatrix} -1 \\ 1 \\ 1 \\ -1 \end{pmatrix}, 0 \right),$$

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

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

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

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 (3). 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 (3) reads as

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

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

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 [19].

Theorem 7.1 ([19]). *Let $k \geqslant 2$ and $p \leqslant \frac{k-1}{k}$. Assume $\mathcal{F} \subset \mathcal{P}([n])$ is k-wise intersecting, that is, for all $F_1, \ldots, F_k \in \mathcal{F}$*

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

Then $\mu_p(\mathcal{F}) \leqslant p$, where μ_p stands the p-biased measure. This bound is attained by dictators, uniquely (unless $k = 2$ and $p = 1/2$).

Proof. We will show that the theorem holds for all $\frac{k-2}{k-1} \leqslant p \leqslant \frac{k-1}{k}$. If $p < \frac{k-2}{k-1}$ then we can deduce the theorem by applying it for $k := k - 1$, since a k-wise intersecting family is also $(k-1)$-wise intersecting.

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

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

from which it follows that its smallest eigenvalue is $\lambda_0 = \frac{\frac{k-2}{k-1}-p}{1-p} \leqslant 0$, hence $\frac{1}{1-\lambda_0} = (k-1)(1-p)$. Furthermore, $\lambda_0 > -1$ as long as $p < 1 - \frac{1}{2(k-1)}$, which is automatically satisfied when $k > 2$.

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 (Theorem 1.4), we conclude

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

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 $X^{\otimes n}$. $\square$

Appendix A. More details on examples from Section 4.1

Example A.1. We start by describing the weighted hypergraph (V, μ). The set of vertices is $V = [3] \times [m]$. If x, y, z are of different color, then $\mu([x, y, z]) = \frac{1}{6m^3}$. The skeleton is the complete tripartite graph $K_{m,m,m}$, and so its eigenvalues are simply the eigenvalues of $K_{m,m,m}$, normalized by the regularity $2m$. The eigenvalues of $K_{m,m,m}$ are readily seen to be $2m, 0, -m$, and so $\lambda_0 = -\frac{1}{2}$. The link of any vertex is the complete bipartite graph $K_{m,m}$, whose eigenvalues are $m, -m$. After normalizing by the regularity m, we get $\lambda_1 = -1$. Altogether, the Link bound is $1 - \frac{2}{3} \cdot \frac{1}{2} = \frac{2}{3}$.

Example A.2. The weighted hypergraph (V, μ) has the vertex set $V = [3] \times [m]$. If x, y, z are distinct vertices of the same color, then $\mu([x, y, z]) = \frac{1}{6m(m-1)(2m-1)}$, and if they are distinct vertices of two different colors, then $\mu([x, y, z]) = \frac{1}{12m(m-1)(2m-1)}$.

We compute $\mu([x, y]) = \frac{1}{3m(2m-1)}$ if $x \neq y$ have the same color and $\mu([x, y]) = \frac{1}{6m(2m-1)}$ if they have different color. Therefore the normalized adjacency operator is

$$T_X = \frac{1}{4m-2}\begin{pmatrix} 2(J-I) & J & J \\ J & 2(J-I) & J \\ J & J & 2(J-I) \end{pmatrix} = \frac{1}{4m-2}\begin{pmatrix} 2 & 1 & 1 \\ 1 & 2 & 1 \\ 1 & 1 & 2 \end{pmatrix} \otimes J - \frac{1}{2m-1}I.$$

The eigenvalues of the small matrix are $4, 1$, and so the eigenvalues of T_X are $1, \frac{m-2}{4m-2}, -\frac{1}{2m-1}$, hence $\lambda_0 = -\frac{1}{2m-1}$.

The link of a vertex x, a graph over $3m-1$ vertices, has the following form. If y, z have the same color, then $\mu_x([y, z]) = \frac{1}{2m-1}$. If y, z have the the same color but not that of x, or if exactly one of them has the color of x, then $\mu_x([y, z]) = \frac{1}{4m-2}$. Therefore the normalized adjacency operator (with blocks of size $m-1, m, m$) is

$$T_{X_x} = \frac{1}{4m-4}\begin{pmatrix} 2(J-I) & J & J \\ 2J & 2(J-I) & 0 \\ 2J & 0 & 2(J-I) \end{pmatrix} = \frac{1}{4m-4}\begin{pmatrix} 2J & J & J \\ 2J & 2J & 0 \\ 2J & 0 & 2J \end{pmatrix} - \frac{1}{2m-2}I.$$

The matrix on the right has rank 2. Its column space is spanned by the eigenvectors $(1, 1, 1)$ and $(0, 1, -1)$, corresponding to the eigenvalues $4m-2$ and $2m$, respectively. Therefore the eigenvalues of T_{X_v} are $1, \frac{1}{2}, -\frac{1}{2m-2}$. In particular, $\lambda_1 = -\frac{1}{2m-2}$.

Altogether, the Link bound is $1 - \frac{2m-1}{2m} \cdot \frac{2m-2}{2m-1} = \frac{1}{m}$.

EXAMPLE A.3. In this example, the weighted hypergraph $X = (V, \mu)$ has vertices $V = [2] \times [m]$. If x, y, z are distinct vertices, not all of the same color, then $\mu([x, y, z]) = \frac{1}{6m^2(m-1)}$.

We compute $\mu([x, y]) = \frac{1}{6m(m-1)}$ if x, y are distinct vertices of the same color, and $\mu([x, y]) = \frac{1}{3m^2}$ if x, y do not have the same color. Therefore the normalized adjacency operator is

$$T_X = \begin{pmatrix} \frac{1}{3m-3}(J-I) & \frac{2}{3m}J \\ \frac{2}{3m}J & \frac{1}{3m-3}(J-I) \end{pmatrix} = \begin{pmatrix} \frac{1}{3m-3} & \frac{2}{3m} \\ \frac{2}{3m} & \frac{1}{3m-3} \end{pmatrix} \otimes J - \frac{1}{3m-3}I.$$

The small matrix has eigenvalues $\frac{3m-2}{3m(m-1)}, \frac{-m+2}{3m(m-1)}$, and so the eigenvalues of T_X are $1, -\frac{1}{3m-3}, -\frac{1}{3}$. This shows that $\lambda_0 = -\frac{1}{3}$.

The link of a vertex x consists of a complete bipartite graph between $m-1$ and m vertices together with a clique on the larger side. Therefore the normalized adjacency operator (with blocks of size $m-1, m$) is

$$T_{X_x} = \begin{pmatrix} 0 & \frac{1}{m}J \\ \frac{1}{2m-2}J & \frac{1}{2m-2}(J-I) \end{pmatrix} = \begin{pmatrix} \frac{1}{2m-2}I & \frac{1}{m}J \\ \frac{1}{2m-2}J & \frac{1}{2m-2}J \end{pmatrix} - \frac{1}{2m-2}I.$$

The matrix on the right has the following eigenvectors: $(1, 1)$ with eigenvalue $\frac{2m-1}{2m-2}$; $(1, -\frac{1}{2})$ with eigenvalue $-\frac{m-2}{2m-2}$; $(x, 0)$ with eigenvalue $\frac{1}{2m-2}$ for any $x \perp 1$ (multiplicity $m-2$); and $(0, y)$ with eigenvalue 0 for any $y \perp 1$ (multiplicity $m-1$). Therefore the eigenvalues of T_{X_x} are $1, -\frac{1}{2}, 0, -\frac{1}{2m-2}$. In particular, $\lambda_1 = -\frac{1}{2}$.

Altogether, the Link bound is $1 - \frac{3}{4} \cdot \frac{2}{3} = \frac{1}{2}$.

EXAMPLE A.4. Here the weighted hypergraph $X = (V, \mu)$ is simply the union of two disjoint 3-uniform m-cliques. That is, $V = [2] \times [m]$, and $\mu([x, y, z]) = \frac{1}{2m(m-1)(m-2)}$ whenever x, y, z are distinct vertices having the same color.

The skeleton is simply the union of two m-cliques, and so its eigenvalues are simply the eigenvalues of K_m, normalized by the regularity $m-1$. In particular, $\lambda_0 = -\frac{1}{m-1}$. Similarly, the link of a vertex is a single $(m-1)$-clique, and so $\lambda_1 = -\frac{1}{m-2}$. It follows that the Link bound is $1 - \frac{m-1}{m} \cdot \frac{m-2}{m-1} = \frac{2}{m}$.

References

[1] Rudolf Ahlswede and Gyula O. H. Katona, *Contributions to the geometry of Hamming spaces*, Discrete Math. **17** (1977), no. 1, 1–22.

[2] Noga Alon, Irit Dinur, Ehud Friedgut, and Benny Sudakov, *Graph products, Fourier analysis and spectral techniques*, Geom. Funct. Anal. **14** (2004), no. 5, 913–940.

[3] Noga Alon and Eyal Lubetzky, *Independent sets in tensor graph powers*, J. Graph Theory **54** (2007), no. 1, 73–87.

[4] Christine Bachoc, Anna Gundert, and Alberto Passuello, *The theta number of simplicial complexes*, Israel J. Math. **232** (2019), no. 1, 443–481.

[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*, J. Combin. Theory Ser. B **133** (2018), 88–121.

[6] Etienne de Klerk and Dmitrii V. Pasechnik, *A note on the stability number of an orthogonality graph*, European J. Combin. **28** (2007), no. 7, 1971–1979.

[7] Phillippe Delsarte, *An algebraic approach to the association schemes of coding theory*, Philips Res. Rep. Suppl. (1973), no. 10, vi+97.

[8] Irit Dinur and Ehud Friedgut, *Intersecting families are essentially contained in juntas*, Combin. Probab. Comput. **18** (2009), no. 1-2, 107–122.

[9] Irit Dinur, Ehud Friedgut, and Oded Regev, *Independent sets in graph powers are almost contained in juntas*, Geom. Funct. Anal. **18** (2008), no. 1, 77–97.

[10] Irit Dinur, Elchanan Mossel, and Oded Regev, *Conditional hardness for approximate coloring*, SIAM J. Comput. **39** (2009), no. 3, 843–873.

[11] Irit Dinur and Samuel Safra, *On the hardness of approximating minimum vertex cover*, Ann. of Math. (2) **162** (2005), no. 1, 439–485.

[12] David Ellis, Yuval Filmus, and Ehud Friedgut, *Triangle-intersecting families of graphs*, J. Eur. Math. Soc. (JEMS) **14** (2012), no. 3, 841–885.

[13] David Ellis, Ehud Friedgut, and Haran Pilpel, *Intersecting families of permutations*, J. Amer. Math. Soc. **24** (2011), no. 3, 649–682.

[14] Paul Erdős, *A problem on independent r-tuples*, Ann. Univ. Sci. Budapest. Eötvös Sect. Math. **8** (1965), 93–95.

[15] Paul Erdős, Chao Ko, and Richard Rado, *Intersection theorems for systems of finite sets*, Quart. J. Math. Oxford Ser. (2) **12** (1961), 313–320.

[16] Yuval Filmus and Ferdinand Ihringer, *Boolean degree 1 functions on some classical association schemes*, J. Combin. Theory Ser. A **162** (2019), 241–270.

[17] Péter Frankl, *On Sperner families satisfying an additional condition*, J. Combinatorial Theory Ser. A **20** (1976), no. 1, 1–11.

[18] _____, *Asymptotic solution of a Turán-type problem*, Graphs Combin. **6** (1990), no. 3, 223–227.

[19] Péter Frankl and Norihide Tokushige, *Weighted multiply intersecting families*, Studia Sci. Math. Hungar. **40** (2003), no. 3, 287–291.

[20] Ehud Friedgut, *Boolean functions with low average sensitivity depend on few coordinates*, Combinatorica **18** (1998), no. 1, 27–35.

[21] _____, *On the measure of intersecting families, uniqueness and stability*, Combinatorica **28** (2008), no. 5, 503–528.

[22] Ehud Friedgut and Oded Regev, *Kneser graphs are like Swiss cheese*, Discrete Anal. (2018), Paper no. 2 (18 pages).

[23] Dion Gijswijt, *Matrix algebras and semidefinite programming techniques for codes*, Ph.D. thesis, University of Amsterdam, 2005.

[24] Dion Gijswijt, Alexander Schrijver, and Hajime Tanaka, *New upper bounds for nonbinary codes based on the Terwilliger algebra and semidefinite programming*, J. Combin. Theory Ser. A **113** (2006), no. 8, 1719–1731.

[25] Chris Godsil and Karen Meagher, *Erdős–Ko–Rado theorems: algebraic approaches*, Cambridge Studies in Advanced Mathematics, vol. 149, Cambridge University Press, Cambridge, 2016.

[26] Konstantin Golubev, *On the chromatic number of a simplicial complex*, Combinatorica **37** (2017), no. 5, 953–964.

[27] Konstantin Golubev and Ori Parzanchevski, *Spectrum and combinatorics of two-dimensional Ramanujan complexes*, Israel J. Math. **230** (2019), no. 2, 583–612.

[28] Anna Gundert and May Szedlák, *Higher dimensional discrete Cheeger inequalities*, J. Comput. Geom. **6** (2015), no. 2, 54–71.

[29] Willem Haemers, *A generalization of the Higman-Sims technique*, Nederl. Akad. Wetensch. Indag. Math. **40** (1978), no. 4, 445–447.

[30] Alan J. Hoffman, *On eigenvalues and colorings of graphs*, in Selected Papers of Alan J. Hoffman — with Commentary, World Scientific, 2003, pp. 407–419.

[31] Jacek Jendrej, Krzysztof Oleszkiewicz, and Jakub O. Wojtaszczyk, *On some extensions of the FKN theorem*, Theory Comput. **11** (2015), 445–469.

[32] Nathan Linial and Yuval Peled, *On the phase transition in random simplicial complexes*, Ann. of Math. (2) **184** (2016), no. 3, 745–773.

[33] László Lovász, *On the Shannon capacity of a graph*, IEEE Trans. Inform. Theory **25** (1979), no. 1, 1–7.

[34] Alexander Lubotzky, *Ramanujan complexes and high dimensional expanders*, Jpn. J. Math. **9** (2014), no. 2, 137–169.

[35] Willem Mantel, *Problem 28 (solution by H. Gouwentak, W. Mantel, J. Teixeira de Mattes, F. Schuh and W. A. Wythoff)*, Wiskundige Opgaven **10** (1907), 60–61.

[36] Robert J. McEliece, Eugene R. Rodemich, Howard Rumsey, Jr., and Lloyd R. Welch, *New upper bounds on the rate of a code via the Delsarte–MacWilliams inequalities*, IEEE Trans. Inform. Theory **IT-23** (1977), no. 2, 157–166.

[37] Roy Meshulam, *On subsets of finite abelian groups with no 3-term arithmetic progressions*, J. Combin. Theory Ser. A **71** (1995), no. 1, 168–172.

[38] Elchanan Mossel, Ryan O'Donnell, and Krzysztof Oleszkiewicz, *Noise stability of functions with low influences: invariance and optimality*, Ann. of Math. (2) **171** (2010), no. 1, 295–341.

[39] Michael Navon and Alex Samorodnitsky, *Linear programming bounds for codes via a covering argument*, Discrete Comput. Geom. **41** (2009), no. 2, 199–207.

[40] Ori Parzanchevski, *Mixing in high-dimensional expanders*, Combin. Probab. Comput. **26** (2017), no. 5, 746–761.

[41] Ori Parzanchevski and Ron Rosenthal, *Simplicial complexes: spectrum, homology and random walks*, Random Structures Algorithms **50** (2017), no. 2, 225–261.

[42] Ori Parzanchevski, Ron Rosenthal, and Ran J. Tessler, *Isoperimetric inequalities in simplicial complexes*, Combinatorica **36** (2016), no. 2, 195–227.

[43] Alexander Schrijver, *A comparison of the Delsarte and Lovász bounds*, IEEE Trans. Inform. Theory **25** (1979), no. 4, 425–429.

[44] _____, *New code upper bounds from the Terwilliger algebra and semidefinite programming*, IEEE Trans. Inform. Theory **51** (2005), no. 8, 2859–2866.

[45] Michael Tait and Josh Tobin, *Three conjectures in extremal spectral graph theory*, J. Combin. Theory Ser. B **126** (2017), 137–161.

[46] Paul M. Weichsel, *The Kronecker product of graphs*, Proc. Amer. Math. Soc. **13** (1962), 47–52.

[47] Richard M. Wilson, *The exact bound in the Erdős–Ko–Rado theorem*, Combinatorica **4** (1984), no. 2-3, 247–257.

[48] Yufei Zhao, *Graph theory and additive combinatorics*, `https://ocw.mit.edu/courses/mathematics/18-217-graph-theory-and-additive-combinatorics-fall-2019/lecture-notes/MIT18_217F19_full_notes.pdf`, 2019.

Yuval Filmus, The Henry and Marilyn Taub Faculty of Computer Science, Technion — Israel Institute of Technology, Haifa, Israel.
E-mail : `yuvalfi@cs.technion.ac.il`

Konstantin Golubev, D-MATH, ETH Zurich, Switzerland. Moscow Center for Fundamental and Applied Mathematics, Russia.
E-mail : `golubevk@ethz.ch`

Noam Lifshitz, Einstein Institute of Mathematics, Hebrew University, Jerusalem, Israel.
E-mail : `noam.lifshitz@gmail.com`