

HYPERGRAPH INDEPENDENT SETS

JONATHAN CUTLER AND A.J. RADCLIFFE

ABSTRACT. The study of extremal problems related to independent sets in hypergraphs is a problem that has generated much interest. There are a variety of types of independent sets in hypergraphs depending on the number of vertices from an independent set allowed in an edge. We say that a subset of vertices is *j-independent* if its intersection with any edge has size strictly less than j . The Kruskal-Katona theorem implies that in an r -uniform hypergraph with a fixed size and order, the hypergraph with the most r -independent sets is the lexicographic hypergraph. In this paper, we use a hypergraph regularity lemma, along with a technique developed by Loh, Pikhurko, and Sudakov [10], to give an asymptotically best possible upper bound on the number of j -independent sets in an r -uniform hypergraph.

1. INTRODUCTION

The study of independent sets in graphs has a long history. Recently, extremal problems related to maximizing the number of independent sets in graphs have been an active area of research. Kahn [7] determined which regular bipartite graphs with n vertices have the most independent sets, and his theorem was recently extended to all regular graphs by Zhao [15]. It is a consequence of the Kruskal-Katona theorem [9, 8] that the lexicographic graph has the largest number of independent sets among graphs of fixed order and size (see, e.g., [2]). Independent sets in hypergraphs have also been well-studied. Much of this research has focused on determining algorithms for finding independent sets in r -uniform hypergraphs (see, e.g., [14]).

In this paper we determine asymptotically the maximum number of independent sets possible for an r -uniform hypergraph on n vertices and m edges. Independent sets in hypergraphs are a bit more complicated than those in graphs since they can be defined in a number of ways depending on how many vertices from an independent set are allowed in an edge.

For us, a *hypergraph* \mathcal{H} is an ordered pair (V, \mathcal{E}) where V is a finite set and $\mathcal{E} \subseteq \mathcal{P}(V)$. We say a hypergraph $\mathcal{H} = (V, \mathcal{E})$ is *r-uniform* if $\mathcal{E} \subseteq \binom{V}{r}$, where $\binom{V}{r} = \{E \subseteq V : |E| = r\}$. For a hypergraph $\mathcal{H} = (V, \mathcal{E})$, we write $n(\mathcal{H})$ for $|V|$ and $e(\mathcal{H})$ for $|\mathcal{E}|$.

Definition 1. For an r -uniform hypergraph $\mathcal{H} = (V, \mathcal{E})$ and an integer j with $1 \leq j \leq r$, we say that a set $I \subseteq V(\mathcal{H})$ is *j-independent* if $|I \cap E| < j$ for all $E \in \mathcal{E}$. Let $\mathcal{I}_j(\mathcal{H})$ be the collection of all j -independent subsets of V and $i_j(\mathcal{H}) = |\mathcal{I}_j(\mathcal{H})|$.

Example 1. If G is a graph, i.e., a 2-uniform hypergraph, then $\mathcal{I}_2(G)$ is the collection of all independent sets in the graph G . On the other hand, $\mathcal{I}_1(G)$ is simply the collection of subsets disjoint from all edges of G . Thus, $i_1(G) = 2^{n_0}$ where n_0 is the number of isolated vertices in G .

In the case $j = r$, the maximum number of independent sets possible in an r -uniform hypergraph is known exactly as a consequence of the Kruskal-Katona theorem. For the sake of completeness, we include this result and its proof.

2000 *Mathematics Subject Classification.* 05C65, 05D05, 05C69.

Key words and phrases. Independent sets, hypergraphs, regularity lemma.

Theorem 1. *If $\mathcal{H} = (V, \mathcal{E})$ is an r -uniform hypergraph with $n(\mathcal{H}) = n$ and $e(\mathcal{H}) = m$, then*

$$i_r(\mathcal{H}) \leq i_r(\mathcal{L}_{n,m}),$$

where $\mathcal{L}_{n,m}$ is the r -graph on $[n] = \{1, 2, \dots, n\}$ whose edges are the first m elements in the lexicographic ordering¹ on $\binom{[n]}{r}$.

Proof. Let us write $\mathcal{I}_r^{(k)}(\mathcal{H})$ for $\mathcal{I}_r(\mathcal{H}) \cap \binom{V}{k}$. Note that $I \in \mathcal{I}_r^{(k)}(\mathcal{H})$ if and only if I is not in the upper shadow of \mathcal{H} on level k , the set $\partial^{(k)}\mathcal{H} = \left\{ B \in \binom{V}{k} : \exists E \text{ an edge of } \mathcal{H} \text{ such that } E \subseteq B \right\}$. Thus

$$\begin{aligned} i_r(\mathcal{H}) &= \sum_{k=0}^n \left| \mathcal{I}_r^{(k)}(\mathcal{H}) \right| \\ &= \sum_{k=0}^{r-1} \binom{n}{k} + \sum_{k=r}^n \left(\binom{n}{k} - \left| \partial^{(k)}\mathcal{H} \right| \right) \\ &\leq \sum_{k=0}^{r-1} \binom{n}{k} + \sum_{k=r}^n \left(\binom{n}{k} - \left| \partial^{(k)}\mathcal{L}_{n,m} \right| \right) \\ &= i_r(\mathcal{L}_{n,m}), \end{aligned}$$

where the inequality $|\partial^{(k)}\mathcal{H}| \geq |\partial^{(k)}\mathcal{L}_{n,m}|$ is the content of the Kruskal-Katona theorem. \square

Maximizing the number of 1-independent sets is trivial, since for an r -uniform hypergraph \mathcal{H} with n vertices,

$$i_1(\mathcal{H}) \leq 2^{n-n^*},$$

where n^* is the unique integer for which $\binom{n^*-1}{r} < e(\mathcal{H}) \leq \binom{n^*}{r}$. Just as in Example 1 for graphs, 1-independent sets in a hypergraph \mathcal{H} are simply subsets of the isolated vertices of \mathcal{H} —vertices not in any edge of \mathcal{H} . This bound is achieved by any hypergraph whose edges are contained in some set of size n^* .

For $1 < j < r$ the problem of maximizing the number of j -independent sets in r -uniform hypergraphs is open. If we follow the approach of Theorem 1 and think about $\mathcal{I}_j^{(k)}(\mathcal{H})$ for some $j < r$, we have

$$\mathcal{I}_j^{(k)}(\mathcal{H}) = \binom{V}{k} \setminus \partial^{(k)}(\partial_{(j)}\mathcal{H}),$$

where the lower shadow $\partial_{(j)}\mathcal{H} = \left\{ B \in \binom{V}{j} : \exists E \text{ an edge of } \mathcal{H} \text{ such that } B \subseteq E \right\}$. The configuration of sets that minimizes the lower shadow, *viz.*, the colex hypergraph, is very different than that which minimizes the upper shadow, *viz.*, the lex hypergraph. It is thus very difficult to get exact results giving upper bounds on the number of independent sets in hypergraphs.

We give an asymptotically best possible bound on $i_j(\mathcal{H})$ in terms of the number of independent sets in a *split hypergraph*, which we define below. Our approach follows that of Loh, Pikhurko, and Sudakov in [10], where they determine asymptotically the maximum number of q -colorings of a graph G with n vertices and m edges. They use Szemerédi's Regularity Lemma [12] in a clever way; given a regular partition of the vertex set and a q -coloring of G , they associate a coloring of the auxiliary graph. Since the regular partition has a bounded number of parts, there are only a constant number of possible auxiliary colorings and one need only consider, asymptotically, those colorings of G corresponding to a fixed auxiliary coloring. This allows them to get good control on the problem of maximizing the number of q -colorings.

¹For $A, B \subseteq [n]$, we say $A <_{\text{lex}} B$ if and only if $\min(A \triangle B) \in A$.

We adapt this approach to prove an asymptotic bound on the number of j -independent sets in an r -uniform hypergraph. We prove that the way to get many j -independent sets is, asymptotically, to have a very large j -independent set, all of whose subsets are, of course, j -independent. Another way to say this is that the asymptotic extremal graphs will be what we call split hypergraphs, which are analogous to split graphs. In adapting the approach of Loh, Pikhurko, and Sudakov we need to use a hypergraph regularity lemma. In Section 2, we discuss both the hypergraph regularity lemma and preliminary results about split hypergraphs, defined below.




FIGURE 1. The types of edges in the split 3-graph $\mathcal{S}_2^{(3)}(A, B)$.

Definition 2. The j -split r -graph with partition (A, B) , denoted $\mathcal{S}_j^{(r)}(A, B)$, is defined as the r -uniform hypergraph with vertex set $V = A \cup B$ and edge set

$$\left\{ E \in \binom{V}{r} : |E \cap A| < j \right\}.$$

When we are not concerned with the identity of the sets A and B , we write $\mathcal{S}_j^{(r)}(k, n - k)$ for a j -split r -graph with $|A| = k$ and $|B| = n - k$. We let $s_j^{(r)}(k, n - k) = e(\mathcal{S}_j^{(r)}(k, n - k))$. See Figure 1 for an example.

We can now state the main result of the paper. In the statement, and for the rest of the paper, all $o(1)$ terms go to zero as n goes to infinity.

Main Theorem. *Let j and $r \neq 1$ be positive integers with $1 \leq j \leq r$. Given $\eta > 0$ and any r -uniform hypergraph \mathcal{H} on n vertices with $\eta \binom{n}{r} \leq e(\mathcal{H}) \leq (1 - \eta) \binom{n}{r}$, if we let k^* be maximal such that*

$$e(\mathcal{H}) \leq s_j^{(r)}(k^*, n - k^*),$$

then

$$i_j(\mathcal{H}) \leq 2^{(1+o(1))k^*}.$$

In Section 2, in addition to proving some preliminary lemmas, we will show that this result is asymptotically best possible by showing that

$$i_j(\mathcal{S}_j^{(r)}(k, n - k)) = (1 + o(1))2^k,$$

for k/n bounded away from 0 and 1. In Section 3, we prove the Main Theorem.

2. PRELIMINARIES

2.1. Hypergraph regularity. There are now highly sophisticated hypergraph regularity lemmas available. (See, e.g., [1, 5, 6, 11, 13].) We, however, need only a simple version which can be found, e.g., in [3] or [4]. In order to state the regularity lemma, we need first to define an ε -regular partition, the structure guaranteed by the regularity lemma. In what follows, we use the standard

notation that if $\mathcal{H} = (V, \mathcal{E})$ is an r -uniform hypergraph and W_1, W_2, \dots, W_r are disjoint subsets of the vertex set, then

$$\mathcal{H}[W_1, W_2, \dots, W_r] = \{E \in \mathcal{E} : |E \cap W_i| = 1, \forall i \in [r]\}.$$

Definition 3. Let $\mathcal{H} = (V, \mathcal{E})$ be an r -uniform hypergraph. Given $\varepsilon > 0$, we say an r -tuple (W_1, W_2, \dots, W_r) of disjoint subsets of V is ε -regular if for all sequences $(S_i)_1^r$ of subsets $S_i \subseteq W_i$ with $|S_i| \geq \varepsilon |W_i|$ for all $i \in [r]$, we have

$$\left| \frac{e(\mathcal{H}[W_1, W_2, \dots, W_r])}{\prod_1^r |W_i|} - \frac{e(\mathcal{H}[S_1, S_2, \dots, S_r])}{\prod_1^r |S_i|} \right| < \varepsilon.$$

We say a partition $\{V_1, \dots, V_t\}$ is an ε -regular partition of \mathcal{H} if

- 1) $|V_1| \leq |V_2| \leq \dots \leq |V_t| \leq |V_1| + 1$, and
- 2) the r -tuple $(V_{i_1}, V_{i_2}, \dots, V_{i_r})$ is ε -regular for all but $\varepsilon \binom{t}{r}$ of the r -sets $\{i_1, i_2, \dots, i_r\}$ in $\binom{[t]}{r}$.

The following hypergraph regularity lemma can be read out of, for example, a result of Czygrinow and Rödl [3].

Theorem 2. For all $r, m \in \mathbb{N}$ and $\varepsilon > 0$, there exists $M, L \in \mathbb{N}$ such that given any r -uniform hypergraph $\mathcal{H} = (V, \mathcal{E})$ with $|V| \geq L$, there is an ε -regular partition $\{V_1, \dots, V_t\}$ of \mathcal{H} with $m \leq t \leq M$.

2.2. Split hypergraphs. We now present some characteristics of the j -split r -graphs, $\mathcal{S}_j^{(r)}(k, n - k)$, which were defined in Section 1. To develop intuition, we briefly discuss the $r = 2$ case.

Example 2. Note that in the case when $r = 2$, only two values for j are of interest, namely 1 and 2. (If $j \geq 3$, then $\mathcal{S}_j^{(2)}(k, n - k)$ is a complete graph.) In the case when $j = 1$, we have that $\mathcal{S}_1^{(2)}(k, n - k)$ is the disjoint union of the empty graph E_k and the complete graph K_{n-k} . When $j = 2$, similarly, $\mathcal{S}_2^{(2)}(k, n - k)$ is the join² of E_k and K_{n-k} .

We now prove a sequence of lemmas concerning hypergraphs which we will need for the proof of the Main Theorem. The first lemma gives the number of edges and j -independent sets in the split r -graphs. For convenience, we write

$$\binom{n}{\leq k} = \sum_{i=0}^k \binom{n}{i}.$$

From this point forward, we will fix an integer r and will discuss r -uniform hypergraphs. Also, we will fix an integer j with $1 \leq j \leq r$. Our aim is to prove the main theorem for these values of r and j , which will appear in our lemmas without further comment.

Lemma 3. Let n be a positive integer with $n \geq r$. If $|A| = k$ and $|B| = n - k$, then the number of edges in $\mathcal{S} = \mathcal{S}_j^{(r)}(A, B)$ is

$$e(\mathcal{S}) = \sum_{i=0}^{j-1} \binom{k}{i} \binom{n-k}{r-i}.$$

If $j \leq k \leq n - r + j - 1$, i.e., \mathcal{S} is neither complete nor empty, then the number of j -independent sets in \mathcal{S} is

$$i_j(\mathcal{S}) = 2^k + \binom{n}{\leq (j-1)} - \binom{k}{\leq (j-1)}.$$

²The join of graphs G and H is the graph with vertex set the disjoint union of $V(G)$ and $V(H)$ and edge set $E(G) \cup E(H) \cup \{xy : x \in V(G), y \in V(H)\}$.

Proof. The first calculation is straightforward. For the second, it is easy to check that in the relevant range of k , any set of j vertices not contained in A is contained in an edge. \square

Corollary 4. *Given $0 < \xi < 1/2$, we have*

$$i_j \left(\mathcal{S}_j^{(r)}(k, n - k) \right) = (1 + o(1))2^k,$$

provided $\xi < k/n < 1 - \xi$. \square

The next technical lemma bounds the difference in the number of edges between split graphs with adjacent values of k . We need to show that the difference is large to ensure that changing k by a constant multiple of n changes the proportion of edges in the hypergraph by a constant amount.

Lemma 5. *Given $0 < \xi < 1/2$ there exists $\zeta > 0$ such that whenever n is an integer with $n \geq 2 \max(j, r - j)/\xi$ and $k \in [n]$ satisfies $\frac{k}{n} \in (\xi, 1 - \xi)$, we have*

$$s_j^{(r)}(k - 1, n - k + 1) \geq s_j^{(r)}(k, n - k) + \zeta n^{r-1}.$$

Proof. Writing $\mathcal{S} = \mathcal{S}_j^{(r)}(k - 1, n - k + 1)$ and $\mathcal{S}' = \mathcal{S}_j^{(r)}(k, n - k)$, we see that \mathcal{S} contains \mathcal{S}' and

$$\begin{aligned} e(\mathcal{S}) - e(\mathcal{S}') &= \binom{k-1}{j-1} \binom{n-k}{r-j} \\ &\geq \frac{(k-j)^{j-1}}{(j-1)!} \cdot \frac{(n-k-r+j)^{r-j}}{(r-j)!} \\ &\geq \frac{1}{(j-1)!(r-j)!} \left(\frac{\xi n}{2} \right)^{j-1} \left(\frac{\xi n}{2} \right)^{r-j} \\ &= \frac{\xi^{r-1}}{2^{r-1}(j-1)!(r-j)!} \cdot n^{r-1}. \end{aligned}$$

We can clearly set $\zeta = \xi^{r-1} / (2^{r-1}(j-1)!(r-j)!)$. \square

Our last lemma discusses the relationship between the number of edges in the split hypergraph $\mathcal{S}_j^{(r)}(k, n - k)$ and the ratio k/n . We make the following definition.

Definition 4. Given an integer $n \geq r$ and e with $0 \leq e \leq \binom{n}{r}$, we write

$$k^*(n, e) = \max \left\{ k \leq n : s_j^{(r)}(k, n - k) \geq e \right\},$$

for the largest k such that the split graph $\mathcal{S}_j^{(r)}(k, n - k)$ has at least e edges. This is clearly well-defined since $s_j^{(r)}(0, n) = \binom{n}{r}$ and $s_j^{(r)}(n, 0) = 0$. Note further that $s_j^{(r)}(k, n - k)$ is a decreasing function of k for fixed n .

We prove that if $e / \binom{n}{r}$ is bounded away from 0 and 1 then so is $k^*(n, e)/n$.

Lemma 6. *Given $\nu \in (0, 1/2)$, there exists $\xi = \xi(\nu) \in (0, 1/2)$ such that for n sufficiently large (as a function of ν), we have the following: if e satisfies*

$$\nu \binom{n}{r} < e < (1 - \nu) \binom{n}{r},$$

then $k^ = k^*(n, e)$ satisfies*

$$\xi < \frac{k^*}{n} < 1 - \xi.$$

Proof. For the lower bound (by Lemma 3), we want to show that there is a ξ such that

$$s_j^{(r)}(\xi n, n - \xi n) = \sum_{i=0}^{j-1} \binom{\xi n}{i} \binom{n - \xi n}{r - i} \geq (1 - \nu) \binom{n}{r}.$$

Pick ξ_1 with $0 < \xi_1 < 1 - (1 - \nu)^{1/r}$. To show the above with $\xi = \xi_1$, we bound the sum by the $i = 0$ term. Thus,

$$\begin{aligned} s_j^{(r)}(\xi_1 n, n - \xi_1 n) &= \sum_{i=0}^{j-1} \binom{\xi_1 n}{i} \binom{n - \xi_1 n}{r - i} \\ &\geq \binom{(1 - \xi_1)n}{r} \\ &= (1 - o(1))(1 - \xi_1)^r \binom{n}{r} \\ &> (1 - \nu) \binom{n}{r}, \end{aligned}$$

for n sufficiently large. In the other direction pick ξ_2 with $0 < \xi_2 < \nu / \left(j \binom{r}{\lfloor r/2 \rfloor} \right)$. We let i' be the value of i with $1 \leq i \leq j - 1$ that maximizes $\binom{(1 - \xi_2)n}{i} \binom{\xi_2 n}{r - i}$. Note that $i' \leq j - 1 < r$. Now,

$$\begin{aligned} s_j^{(r)}(n - \xi_2 n, \xi_2 n) &= \sum_{i=0}^{j-1} \binom{n - \xi_2 n}{i} \binom{\xi_2 n}{r - i} \\ &\leq j \binom{(1 - \xi_2)n}{i'} \binom{\xi_2 n}{r - i'} \\ &< j (1 - \xi_2)^{i'} \xi_2^{r - i'} \binom{n}{i'} \binom{n}{r - i'}. \end{aligned}$$

Since we want an upper bound, we can neglect the factors of $1 - \xi_2$; since $i' < r$, there is at least one factor of ξ_2 . So,

$$\begin{aligned} s_j^{(r)}(n - \xi_2 n, \xi_2 n) &\leq j \xi_2 \binom{n}{i'} \binom{n}{r - i'} \\ &= j \xi_2 (1 + o(1)) \binom{n}{i'} \binom{n - i'}{r - i'} \\ &= j \xi_2 (1 + o(1)) \binom{n}{r} \binom{r}{i'} \\ &\leq j \xi_2 (1 + o(1)) \binom{n}{r} \binom{r}{\lfloor r/2 \rfloor} \\ &< \nu \binom{n}{r}, \end{aligned}$$

for n sufficiently large. Letting $\xi = \min(\xi_1, \xi_2)$, the lemma follows. \square

3. PROOF OF MAIN THEOREM

We restate the main theorem before presenting its proof.

Main Theorem. Let j and $r \neq 1$ be positive integers with $1 \leq j \leq r$. Given $\eta > 0$ and any r -uniform hypergraph \mathcal{H} on n vertices with $\eta \binom{n}{r} \leq e(\mathcal{H}) \leq (1 - \eta) \binom{n}{r}$, if we let k^* be maximal such that

$$e(\mathcal{H}) \leq s_j^{(r)}(k^*, n - k^*),$$

then

$$i_j(\mathcal{H}) \leq 2^{(1+o(1))k^*}.$$

Proof. Given $\epsilon > 0$, we want to show that for n sufficiently large and \mathcal{H} an r -uniform hypergraph with $\eta \binom{n}{r} \leq e(\mathcal{H}) \leq (1 - \eta) \binom{n}{r}$, we have

$$(1) \quad i_j(\mathcal{H}) \leq 2^{(1+\epsilon)k^*(n,e)},$$

where we have written e for $e(\mathcal{H})$. We will proceed by introducing a parameter δ which governs our use of the regularity lemma; at the end of the proof we will choose δ sufficiently small so as to achieve the bound in (1). Our proof proceeds in a sequence of steps.

Step 1. Given $0 < \delta < 1$, there exists a “cleaned-up” subhypergraph \mathcal{H}' of \mathcal{H} and a δ -regular partition $\{V_1, V_2, \dots, V_t\}$ of \mathcal{H} such that

- a) $e(\mathcal{H}') \geq e(\mathcal{H}) - \delta n^r$,
- b) all edges of \mathcal{H}' span r parts of the partition, and
- c) all subgraphs $\mathcal{H}'[V_{i_1}, V_{i_2}, \dots, V_{i_r}]$ are either empty or δ -regular with density at least δ .

To show this, first apply Theorem 2 with $\varepsilon = \delta/4$ and m sufficiently large that

$$\prod_{i=1}^{r-1} \left(1 - \frac{i}{m}\right) > 1 - \frac{\delta}{2}$$

to get a suitable partition $\{V_1, V_2, \dots, V_t\}$. We get \mathcal{H}' by first deleting all edges that do not span r parts of the partition. Taking into account the fact that the sizes of the V_i are each within one of n/t , and the number of parts in the partition is bounded by M , the number of r -sets in $V(\mathcal{H})$ that do not span r parts of the partition is at most

$$\begin{aligned} (1 + o(1)) \left[\binom{n}{r} - \frac{n^r \prod_{i=1}^{r-1} \left(1 - \frac{i}{t}\right)}{r!} \right] &= (1 + o(1)) \binom{n}{r} \left[1 - \frac{n^r}{(n)_r} \prod_{i=1}^{r-1} \left(1 - \frac{i}{t}\right) \right] \\ &\leq (1 + o(1)) \binom{n}{r} \left[1 - \prod_{i=1}^{r-1} \left(1 - \frac{i}{m}\right) \right] \\ &< (1 + o(1)) \frac{\delta}{2} \binom{n}{r} \\ &\leq (1 + o(1)) \frac{\delta}{4} n^r \\ &< \frac{3\delta}{8} n^r \end{aligned}$$

for n sufficiently large. Certainly, then, the number of edges of \mathcal{H} that do not span r parts of the partition is at most $\frac{3\delta}{8} n^r$. Now, if $\{i_1, i_2, \dots, i_r\}$ has the property that $\mathcal{H}[V_{i_1}, V_{i_2}, \dots, V_{i_r}]$ either has density less than δ or is not δ -regular, we delete all edges of $\mathcal{H}[V_{i_1}, V_{i_2}, \dots, V_{i_r}]$. The total number of such deleted edges is at most

$$\delta \binom{t}{r} \binom{n}{t}^r + \frac{\delta}{4} \binom{t}{r} \binom{n}{t}^r < \frac{5\delta}{8} n^r.$$

(Recall that the partition was chosen to be $(\delta/4)$ -regular.) The hypergraph \mathcal{H}' is the result of deleting all of the above edges from \mathcal{H} .

Step 2. Now we focus our attention on j -independent sets with the property that if they intersect a part V_i , then they have reasonably large intersection with that part.

Definition 5. Given a j -independent set I in \mathcal{H}' we define a subset $\mathcal{D}(I) \subseteq [t]$ by

$$\mathcal{D}(I) = \{i \in [t] : |I \cap V_i| \geq \delta |V_i|\}.$$

We call a j -independent set *robust* if $I \cap V_i = \emptyset$ for all $i \notin \mathcal{D}(I)$. We let \mathcal{R} be the set of robust j -independent sets in \mathcal{H}' . Also, for convenience we define V_D , for $D \subseteq [t]$, to be

$$V_D = \bigcup_{i \in D} V_i.$$

The following lemma proves that a reasonable proportion of j -independent sets are robust.

Lemma 7. *There exists $c_\delta > 0$ such that*

$$|\mathcal{R}| \geq \exp(-c_\delta n) i_j(\mathcal{H}')$$

and c_δ tends to zero as δ tends to zero.

Proof. Define a map $f : \mathcal{I}_j(\mathcal{H}') \rightarrow \mathcal{R}$ by

$$f(I) = I \cap V_{\mathcal{D}(I)}.$$

We show that each $I_0 \in \mathcal{R}$ has at most $\exp((2\delta - \delta \ln \delta)n)$ preimages under f . First note that $|I \setminus f(I)| \leq \delta n$, therefore the number of I such that $f(I) = I_0$ is at most, using Stirling's Formula to bound the largest term,

$$\begin{aligned} \sum_{s=0}^{\delta n} \binom{n}{s} &\leq (1 + \delta n) \exp(\delta n) \left(\frac{1}{\delta}\right)^{\delta n} \\ &\leq \exp(\delta n) \exp(\delta n) \left(\frac{1}{\delta}\right)^{\delta n} = \exp((2\delta - \delta \ln \delta)n). \end{aligned}$$

Setting $c_\delta = 2\delta - \delta \ln \delta$ proves the lemma. \square

Step 3. We have shown that there aren't too many fewer robust j -independent sets than j -independent sets. Now we show that the existence of many robust j -independent sets constrains \mathcal{H}' to be a subgraph of a j -split hypergraph. To be precise, for $D \subseteq [t]$, we let \mathcal{S}_D be the j -split hypergraph $\mathcal{S}_j^{(r)}(V_D, V_D^c)$, and show the following.

Lemma 8. *If $I \in \mathcal{R}$ and $\mathcal{D}(I) = D$, then $\mathcal{H}' \subseteq \mathcal{S}_D$, or, in other words, V_D is j -independent.*

Proof. We need to show that if $\{V_{i_1}, V_{i_2}, \dots, V_{i_r}\}$ is a set of blocks with $|\{i_1, i_2, \dots, i_r\} \cap D| \geq j$, then $\mathcal{H}'[V_{i_1}, V_{i_2}, \dots, V_{i_r}]$ is empty. If it is not empty then it is δ -regular. For $i \in \{i_1, i_2, \dots, i_r\} \cap D$, we have $|I \cap V_i| \geq \delta |V_i|$, so, by the δ -regularity, there is an edge E of \mathcal{H}' with $E \cap V_i \subseteq I \cap V_i$ for all $i \in \{i_1, i_2, \dots, i_r\} \cap D$. In particular, $|E \cap I| \geq j$, contradicting the j -independence of I . \square

Step 4. Finally, we prove (1), that

$$\log_2(i_j(\mathcal{H})) \leq (1 + \epsilon)k^*(n, e).$$

For $D \subseteq [t]$, we let \mathcal{R}_D be the collection of robust j -independent sets supported on D , i.e., $\mathcal{R}_D = \{I \in \mathcal{R} : \mathcal{D}(I) = D\}$. Fix a D^* with $|\mathcal{R}_{D^*}|$ maximal, so that

$$(\dagger) \quad |\mathcal{R}_{D^*}| \geq 2^{-t} |\mathcal{R}| \geq 2^{-M} |\mathcal{R}|.$$

We estimate as follows.

$$\begin{aligned} i_j(\mathcal{H}) &\leq i_j(\mathcal{H}') \\ &\leq \exp(c_\delta n) |\mathcal{R}| \\ &\leq \exp(c_\delta n) 2^M |\mathcal{R}_{D^*}| \\ &\leq \exp(c_\delta n) 2^M 2^{|V_{D^*}|}. \end{aligned}$$

The first inequality follows from the fact that $\mathcal{H}' \subseteq \mathcal{H}$, the second from Lemma 7, and the third from (\dagger) . The final estimate is simply the fact that if I is a robust independent set with $\mathcal{D}(I) = D$, then $I \subseteq V_D$. Taking logarithms, we have

$$(2) \quad \log_2(i_j(\mathcal{H})) \leq \frac{c_\delta n}{\ln 2} + |V_{D^*}| + M.$$

We now need to bound $|V_{D^*}|$.

Lemma 9. *There exists $c'_\delta > 0$ such that*

$$|V_{D^*}| \leq k^*(n, e) + c'_\delta n$$

and c'_δ tends to zero as δ tends to zero.

Proof. Let $e' = e(\mathcal{H}')$ and $k' = k^*(n, e') = \max\{k : s_j^{(r)}(k, n-k) \geq e'\}$. Since \mathcal{H}' is a subgraph of \mathcal{S}_{D^*} , by Lemma 8, we have $|V_{D^*}| \leq k'$. We apply Lemma 6 with $\nu = \eta/2$ to get a suitable $\xi \in (0, 1/2)$. Recall that $e' > \eta \binom{n}{r} - \delta n^r$. So, for n sufficiently large and $\delta < \eta/(4r!)$, we have

$$\frac{\eta}{2} \binom{n}{r} < (\eta - 2r!\delta) \binom{n}{r} \leq e' \leq e(\mathcal{H}) \leq (1 - \eta) \binom{n}{r} < \left(1 - \frac{\eta}{2}\right) \binom{n}{r}.$$

So, Lemma 6 implies that

$$\xi < \frac{k^*(n, e)}{n} \leq \frac{k'}{n} < 1 - \xi.$$

Thus, by Lemma 5, there exists $\zeta > 0$ such that for all k with $k^*(n, e) \leq k \leq k'$, we have

$$s_j^{(r)}(k-1, n-k+1) \geq s_j^{(r)}(k, n-k) + \zeta n^{r-1}.$$

In particular,

$$\begin{aligned} \delta n^r &\geq e - e' \\ &\geq s_j^{(r)}(k^*(n, e) + 1, n - k^*(n, e) - 1) - s_j^{(r)}(k', n - k') \\ &\geq (k' - k^*(n, e) - 1) \zeta n^{r-1}. \end{aligned}$$

Therefore, $k' - k^*(n, e) \leq \frac{\delta n}{\zeta} + 1 \leq \frac{2\delta n}{\zeta}$. We have

$$|V_{D^*}| \leq k' \leq k^*(n, e) + \frac{2\delta}{\zeta} n.$$

We set $c'_\delta = \frac{2\delta}{\zeta}$ and the lemma is proved noting that ζ does not depend on δ . \square

Applying Lemma 9 to (2), we have, for n sufficiently large,

$$\begin{aligned} \log_2(i_j(\mathcal{H})) &\leq \frac{c_\delta n}{\ln 2} + k^*(n, e) + c'_\delta n + M \\ &\leq k^*(n, e) + \left(\frac{c_\delta}{\ln 2} + c'_\delta\right) \xi k^*(n, e) + M \\ &\leq (1 + \epsilon)k^*(n, e). \end{aligned}$$

The penultimate step follows from Lemma 6. The final bound comes from choosing δ sufficiently small. \square

REFERENCES

1. Fan R. K. Chung, *Regularity lemmas for hypergraphs and quasi-randomness*, Random Structures Algorithms **2** (1991), no. 2, 241–252. MR MR1099803 (92d:05117)
2. Jonathan Cutler and A. J. Radcliffe, *Extremal problems for independent set enumeration*, Electron. J. Combin. **18** (2011), no. 1, Research Paper 169, 17.
3. Andrzej Czygrinow and Vojtech Rödl, *An algorithmic regularity lemma for hypergraphs*, SIAM J. Comput. **30** (2000), no. 4, 1041–1066 (electronic). MR MR1786751 (2001j:05088)
4. Alan Frieze and Ravi Kannan, *Quick approximation to matrices and applications*, Combinatorica **19** (1999), no. 2, 175–220. MR MR1723039 (2001i:68066)
5. W. T. Gowers, *Quasirandomness, counting and regularity for 3-uniform hypergraphs*, Combin. Probab. Comput. **15** (2006), no. 1-2, 143–184. MR MR2195580 (2008b:05175)
6. ———, *Hypergraph regularity and the multidimensional Szemerédi theorem*, Ann. of Math. (2) **166** (2007), no. 3, 897–946. MR MR2373376 (2009d:05250)
7. Jeff Kahn, *An entropy approach to the hard-core model on bipartite graphs*, Combin. Probab. Comput. **10** (2001), no. 3, 219–237. MR MR1841642 (2003a:05111)
8. G. Katona, *A theorem of finite sets*, Theory of graphs (Proc. Colloq., Tihany, 1966), Academic Press, New York, 1968, pp. 187–207. MR MR0290982 (45 #76)
9. Joseph B. Kruskal, *The number of simplices in a complex*, Mathematical optimization techniques, Univ. of California Press, Berkeley, Calif., 1963, pp. 251–278. MR MR0154827 (27 #4771)
10. Po-Shen Loh, Oleg Pikhurko, and Benny Sudakov, *Maximizing the number of q -colorings*, to appear in Proc. London Math. Soc.
11. V. Rödl, B. Nagle, J. Skokan, M. Schacht, and Y. Kohayakawa, *The hypergraph regularity method and its applications*, Proc. Natl. Acad. Sci. USA **102** (2005), no. 23, 8109–8113 (electronic). MR MR2167756 (2006g:05095)
12. Endre Szemerédi, *Regular partitions of graphs*, Problèmes combinatoires et théorie des graphes (Colloq. Internat. CNRS, Univ. Orsay, Orsay, 1976), Colloq. Internat. CNRS, vol. 260, CNRS, Paris, 1978, pp. 399–401. MR MR540024 (81i:05095)
13. Terence Tao, *A variant of the hypergraph removal lemma*, J. Combin. Theory Ser. A **113** (2006), no. 7, 1257–1280. MR MR2259060 (2007k:05098)
14. Raphael Yuster, *Finding and counting cliques and independent sets in r -uniform hypergraphs*, Inform. Process. Lett. **99** (2006), no. 4, 130–134. MR MR2236810 (2007a:05092)
15. Yufei Zhao, *The number of independent sets in a regular graph*, Combin. Probab. Comput. **19** (2010), no. 2, 315–320. MR MR2593625

DEPARTMENT OF MATHEMATICAL SCIENCES, MONTCLAIR STATE UNIVERSITY, MONTCLAIR NJ 07043

E-mail address: jonathan.cutler@montclair.edu

DEPARTMENT OF MATHEMATICS, UNIVERSITY OF NEBRASKA-LINCOLN, LINCOLN NE 68588-0130

E-mail address: aradcliffe1@math.unl.edu