

Pattern Avoidance Over a Hypergraph

Benjamin Gunby

Department of Mathematics
Rutgers University, U.S.A.

bg570@connect.rutgers.edu

Maxwell Fishelson

Department of Computer Science
Massachusetts Institute of Technology, U.S.A.

maxfish@mit.edu

Submitted: Sep 16, 2019; Accepted: Nov 12, 2021; Published: Dec 17, 2021

© The authors. Released under the CC BY-ND license (International 4.0).

Abstract

A classic result of Marcus and Tardos (previously known as the Stanley-Wilf conjecture) bounds from above the number of n -permutations ($\sigma \in S_n$) that do not contain a specific sub-permutation. In particular, it states that for any fixed permutation π , the number of n -permutations that avoid π is at most exponential in n . In this paper, we generalize this result. We bound the number of avoidant n -permutations even if they only have to avoid π at specific indices. We consider a k -uniform hypergraph Λ on n vertices and count the n -permutations that avoid π at the indices corresponding to the edges of Λ . We analyze both the random and deterministic hypergraph cases. This problem was originally proposed by Asaf Ferber.

When Λ is a random hypergraph with edge density α , we show that the expected number of Λ -avoiding n -permutations is bounded (both upper and lower) as $\exp(O(n))\alpha^{-\frac{n}{k-1}}$, using a supersaturation version of Füredi-Hajnal.

In the deterministic case we show that, for Λ containing many size L cliques, the number of Λ -avoiding n -permutations is $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$, giving a nontrivial bound with L polynomial in n . Our main tool in the analysis of this deterministic case is the new and revolutionary hypergraph containers method, developed in [1] and [8].

Mathematics Subject Classifications: 05A05, 05A16, 05C65

1 Introduction

Formally, the notion of pattern avoidance is defined as follows.

Definition 1.1. An n -permutation σ **contains** a k -permutation π iff there exist integers $1 \leq x_1 < x_2 < \cdots < x_k \leq n$ such that

$$\pi(i) < \pi(j) \Leftrightarrow \sigma(x_i) < \sigma(x_j)$$

for all i, j . Otherwise, we say σ **avoids** π .

In the late 1980s/early 1990s, Richard P. Stanley and Herbert Wilf independently conjectured that for every permutation π , there exists a constant c_π such that the number of n -permutations avoiding $\pi < c_\pi^n$ for all n . As there are $n! = e^{(1-o(1))n \log n}$ permutations, this bound is non-trivial. This conjecture was later proven by Marcus and Tardos [6] in 2003.

To generalize this result, we first generalize our notion of pattern avoidance in order to account for avoidance only at specific index sets.

Definition 1.2. Let Λ be a k -uniform hypergraph on vertex set $\{1, 2, \dots, n\}$. We say an n -permutation σ **Λ -contains** a k -permutation π iff there exist integers $1 \leq x_1 < x_2 < \dots < x_k \leq n$ such that

$$\pi(i) < \pi(j) \Leftrightarrow \sigma(x_i) < \sigma(x_j)$$

for all i, j **AND** $\{x_1, \dots, x_k\} \in E(\Lambda)$. Otherwise, we say σ **Λ -avoids** π .

In this paper, we analyze the generalized Λ -avoidance problem for both random hypergraphs and fixed hypergraphs, a problem originally posed by Asaf Ferber [2]. When Λ is a random hypergraph with edge density α , we show that, for every permutation fixed $k \in \mathbb{Z}^+$ and $\pi \in S_k$, the number of Λ -avoiding n -permutations is $\exp(O(n))\alpha^{-\frac{n}{k-1}}$ in expectation. We also show that, for fixed Λ , the number of n -permutations Λ -avoiding π is $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$ for all $\epsilon > 0$, as long as Λ is k -uniform and satisfies the following:

Λ contains a collection of L -vertex cliques where each of the n vertices belongs to at least $\delta(\Lambda) \geq 1$ cliques in the collection and at most $\Delta(\Lambda) = O(1)$.

We see that, for $L = n^{\Omega(1)}$, this bound is a non-negligible improvement on the $n^{(1-o(1))n}$ total n -permutations.

A few years after the proposal of Stanley-Wilf, in 1992, Zoltán Füredi and Péter Hajnal proposed a similar conjecture [4] that extended the notion of pattern-avoiding permutations to pattern-avoiding matrices. Essentially, a 0-1 matrix A of size $n \times n$ contains a 0-1 matrix P of size $k \times k$ if there exists a $k \times k$ submatrix of A that has 1-entries at all the locations where P has 1-entries. Formally,

Definition 1.3. For a 0-1 matrix A of size $n \times n$ and a $k \times k$ 0-1 matrix P , we say that A **contains** P iff there exist row indices $1 \leq x_1 < x_2 < \dots < x_k \leq n$ and column indices $1 \leq y_1 < y_2 < \dots < y_k \leq n$ such that

$$P_{ij} = 1 \Rightarrow A_{x_i y_j} = 1$$

for all i, j . Otherwise, we say A **avoids** P . We note that, for A to contain P , we don't require that P be a submatrix of A , but that the 1-entries of P be present in a submatrix of A .

The Füredi-Hajnal conjecture states that, if a 0-1 matrix A of size $n \times n$ avoids a permutation matrix P_π , it has $< c_P n$ 1-entries for some constant c_P in terms of π . Progress was first made on these conjectures by Martin Klazar in 2000 [5], who showed

that the Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture. Then, in 2004, Adam Marcus and Gábor Tardos proved the Füredi-Hajnal conjecture [6]. Combined with Klazar's arguments, a proof of the Stanley-Wilf conjecture was finally achieved.

This notion of pattern-avoiding matrices parallels that of pattern-avoiding permutations, as a permutation σ contains a permutation π if and only if the permutation matrix P_σ contains the permutation matrix P_π . The notion of Λ -avoidance can also be extended to this matrix context, where A must only avoid P on submatrices whose columns correspond to an edge in Λ . Viewing pattern avoidance in this matrix context was the key to proving the Stanley-Wilf conjecture and will be one of the main insights in our analysis.

2 Main Results

When Λ is a random hypergraph, we will prove the following bound.

Theorem 2.1. *Let $k \in \mathbb{Z}$ with $k > 1$, and take $\pi \in S_k$. Then there is some constant $C = C(\pi)$ such that if Λ is the k -uniform Erdős-Rényi random hypergraph on n vertices with edge probability $\alpha \in (0, 1]$, then the expected number of $\sigma \in S_n$ that Λ -avoid π is at most*

$$\exp(Cn)\alpha^{-\frac{n}{k-1}}.$$

Furthermore, for $\alpha \geq n/\binom{n}{k}$, this bound is sharp to within an exponential factor; that is, up to a modification in C (making it potentially negative).

Remark 2.2. When $\alpha = n/\binom{n}{k}$, note that this theorem gives a lower bound of $\exp(Cn) \cdot n!$. Thus $n!$ is correct to an exponential factor when $\alpha < n/\binom{n}{k}$ (the answer can only increase when α decreases). Thus, we have successfully managed to bound the expectation within an exponential for all values of α .

As we will see in Section 3, due to linearity of expectation, Theorem 2.1 reduces to bounding the number of permutations containing few copies of π , for which we will require bounds on the maximal number of ones in a 0-1 matrices containing few copies of the permutation matrix A_π . Both of these bounds may be of independent interest as they give sharp approximations up to respectively an exponential and a constant.

Theorem 2.3. *Let $k \in \mathbb{Z}^+$, $\pi \in S_k$, and let A_π be the $k \times k$ permutation matrix corresponding to π . There exist constants $C = C(\pi)$ and $C' = C'(\pi) > 0$ such that if M is a 0-1 matrix of size $n \times n$ containing a ones, with $Cn \leq a \leq n^2$, then M contains at least $C' \frac{a^{2k-1}}{n^{2k-2}}$ copies of A_π . Furthermore, for a in the given range, this bound is sharp to within a constant factor (depending on π), in the sense that for any a one can always find an M that attains this lower bound to within a constant factor.*

Remark 2.4. This theorem can be thought of as a 'supersaturation' version of Füredi-Hajnal. Indeed, while Füredi-Hajnal states that with $> Cn$ ones at least one copy of A_π is forced, Theorem 2.3 gives a bound on the number of copies of A_π that are forced by any number of ones. We will see in Section 4 how to prove Theorem 2.3 by bootstrapping Füredi-Hajnal.

In a similar but more complicated way to the deduction of Stanley-Wilf from Füredi-Hajnal, we will be able to show the following.

Theorem 2.5. *Let $k \in \mathbb{Z}^+$, $k > 1$ and $\pi \in S_k$. There exists some constant $C = C(\pi)$ and $c = c(\pi)$, $C > c \in \mathbb{R}$ such that for all $m, n \in \mathbb{Z}^{\geq 0}$ with $m \leq \binom{n}{k}$, letting $S_n(m, \pi)$ be the number of permutations in S_n containing at most m copies of π , we have that*

$$\exp(cn) \cdot \max \left(1, \left(\frac{m}{n} \right)^{\frac{n}{k-1}} \right) \leq S_n(m, \pi) \leq \exp(Cn) \cdot \max \left(1, \left(\frac{m}{n} \right)^{\frac{n}{k-1}} \right).$$

Remark 2.6. Notice that c is potentially negative, and thus the lower bound in Theorem 2.5 is only nontrivial if $m > n$. Further note that the theorem simply reduces to Stanley-Wilf when $m = 0$.

In Section 3, we will make the easy deduction of Theorem 2.1 as a corollary of Theorem 2.5. In Section 4, we will prove Theorem 2.3 by bootstrapping Füredi-Hajnal. Finally, in Section 5 we will prove Theorem 2.5.

Remark 2.7. It is natural to ask whether in addition to the expectation result in Theorem 2.1, one can also derive a concentration result on the number of σ that Λ -avoid π . However, this at least does not seem to us to be obvious. For example, if one tries computing the variance of the random variable (in order to for example apply Chebyshev's inequality), one immediately runs into difficulties, which we will now describe.

Notice that the expected number of σ that Λ -avoid π is a sum of indicator random variables, one for each $\sigma \in S_n$, which are 1 if and only if that σ Λ -avoids π . In a variance computation, we would need to compute the covariance of these indicators for σ and σ' . This computation involves simultaneously keeping track of the number of copies of π in σ and σ' , along with the number of index sets (of size k) at which σ and σ' simultaneously contain a copy of π . The interaction of these three quantities seems difficult to deal with.

Of course, Markov's inequality combined with Theorem 2.1 shows that an upper bound of the type given in Theorem 2.1 holds with probability $1 - e^{-n}$. However, since this bound potentially involves changing the constant C , it is quite weak, and given the difficulties above it seems like proving stronger bounds will likely require new ideas.

We will also consider the case when Λ is a fixed graph with particular structure. In particular, we will show the following.

Theorem 2.8. *For every permutation π and any $\epsilon > 0$, the number of n -permutations Λ -avoiding π is $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$ as long as Λ is k -uniform and satisfies the following:*

Λ contains a collection of L -vertex cliques where each of the n vertices belongs to at least $\delta(\Lambda) \geq 1$ cliques in the collection and at most $\Delta(\Lambda) = O(1)$.

In Sections 6 to 10, we will prove Theorem 2.8. The main tool in our analysis will be the hypergraph containers method. The containers method enables us to distribute the vertices of a hypergraph into containers such that every independent set in the hypergraph

belongs to one of the containers. We can apply this method recursively, breaking each container down further into more containers in a branching fashion, to bound the total number of independent sets in a hypergraph.

We will set up a hypergraph whose vertices represent the 1-entries in a matrix and whose edges represent the entries in a submatrix containing P_π with columns $\in E(\Lambda)$. In this context, independent sets correspond to Λ -avoiding matrices. Using the hypergraph containers method, we bound the number of permutation-matrix independent sets, utilizing Füredi-Hajnal to show that the conditions needed to apply the method hold.

In Section 6 we introduce this fixed Λ case. We motivate the constraint that Λ contains cliques of size $\text{poly}(n)$ by demonstrating that Λ with maximal clique of constant size can contain $\Theta(n^k)$ edges and still be avoided by almost all n -permutations. In Section 7, we establish the matrix/hypergraph formulation of the problem. In Section 8, we formally introduce the hypergraph containers lemma and investigate the necessary conditions to apply the lemma in a recursive branching fashion. In Sections 9 and 10, we verify that these conditions are met using two additional lemmas, completing the proof of Theorem 2.8.

Many of the arguments in these sections parallel those presented in a paper [3] by Asaf Ferber, Gweneth Anne McKinley, and Wojciech Samotij. Additionally, the application of the hypergraph container lemma in a recursive branching fashion is adopted from a paper [7] by Morris and Saxton.

Lastly, in Section 11, we will compare Theorems 2.1 and 2.8 and summarize our results.

3 Linearity of Expectation

Fix an integer $k > 1$ and suppose Λ is a random k -uniform hypergraph on $[n]$ with each edge chosen independently at random with edge probability $\alpha \in (0, 1]$. In this case, we may simplify the problem by making use of linearity of expectation. In particular, let us define

$$Av_{n,\Lambda}(\pi) := \{\sigma \in S_n : \sigma \text{ } \Lambda\text{-avoids } \pi\}.$$

Then by linearity of expectation, we have that

$$\mathbb{E}_\Lambda[|Av_{n,\Lambda}(\pi)|] = \sum_{\sigma \in S_n} \Pr[\sigma \text{ } \Lambda\text{-avoids } \pi].$$

This latter probability is simply the probability that none of the copies of π in σ correspond to edges of Λ , which is $(1 - \alpha)^{\# \text{ of copies of } \pi \text{ in } \sigma}$. Therefore,

$$\mathbb{E}_\Lambda[|Av_{n,\Lambda}(\pi)|] = \sum_{\sigma \in S_n} (1 - \alpha)^{\# \text{ of copies of } \pi \text{ in } \sigma}. \tag{1}$$

Thus bounds on the number of permutations containing few copies of π , as given in Theorem 2.5, will give us bounds on our desired quantity $\mathbb{E}_\Lambda[|Av_{n,\Lambda}(\pi)|]$. We now make this argument rigorous.

Deduction of Theorem 2.1 from Theorem 2.5. We first prove the upper bound. By (1),

$$\begin{aligned} \mathbb{E}_\Lambda[|Av_{n,\Lambda}(\pi)|] &= \sum_{\sigma \in S_n} (1 - \alpha)^{\# \text{ of copies of } \pi \text{ in } \sigma} \\ &\leq \sum_{m=0}^{\binom{n}{k}} (1 - \alpha)^m \cdot |\{\sigma \in S_n : \sigma \text{ contains at most } m \text{ copies of } \pi\}|. \end{aligned}$$

By Theorem 2.5, there exists $C = C(\pi)$ such that this is at most

$$\begin{aligned} \sum_{m=0}^{\binom{n}{k}} (1 - \alpha)^m \exp(Cn) \cdot \max\left(1, \left(\frac{m}{n}\right)^{\frac{n}{k-1}}\right) &= \sum_{m=0}^n (1 - \alpha)^m \exp(Cn) \\ &\quad + \sum_{m=n+1}^{\binom{n}{k}} (1 - \alpha)^m \exp(Cn) \left(\frac{m}{n}\right)^{\frac{n}{k-1}} \\ &\leq (n + 1) \exp(Cn) + \frac{\exp(Cn)}{n^{\frac{n}{k-1}}} \sum_{m=n+1}^{\binom{n}{k}} (1 - \alpha)^m m^{\frac{n}{k-1}} \\ &\leq (n + 1) \exp(Cn) + \binom{n}{k} \frac{\exp(Cn)}{n^{\frac{n}{k-1}}} \cdot \max_{m \in \mathbb{R}^+} (1 - \alpha)^m m^{\frac{n}{k-1}} \\ &\leq \exp((C + k)n) \left(1 + n^{-\frac{n}{k-1}} \cdot \max_{m \in \mathbb{R}^+} (1 - \alpha)^m m^{\frac{n}{k-1}}\right) \\ &\leq \exp((C + k)n) \left(1 + n^{-\frac{n}{k-1}} \cdot \max_{m \in \mathbb{R}^+} e^{-\alpha m} m^{\frac{n}{k-1}}\right), \end{aligned}$$

where we are simply bounding our sum by its number of terms times its maximum term, and using the trivial bounds $\binom{n}{k} \leq n^k < (e^n)^k = e^{kn}$, $n + 1 \leq kn < e^{kn}$ for $n \geq 1$ and $k \geq 2$, and $1 - \alpha \leq e^{-\alpha}$. Now, by taking the logarithm and differentiating with respect to m , we see that $e^{-\alpha m} m^{\frac{n}{k-1}}$ is maximized when $-\alpha + \frac{n}{(k-1)m} = 0$, or rearranging, $m = \frac{n}{(k-1)\alpha}$. Substituting, we have that

$$\max_{m \in \mathbb{R}^+} e^{-\alpha m} m^{\frac{n}{k-1}} = \left(\frac{n}{e(k-1)\alpha}\right)^{\frac{n}{k-1}}.$$

Putting this into the calculation from earlier,

$$\begin{aligned} \mathbb{E}_\Lambda[|Av_{n,\Lambda}(\pi)|] &\leq \exp((C + k)n) \left(1 + n^{-\frac{n}{k-1}} \left(\frac{n}{e(k-1)\alpha}\right)^{\frac{n}{k-1}}\right) \\ &= \exp((C + k)n) \left(1 + \left(\frac{1}{e(k-1)\alpha}\right)^{\frac{n}{k-1}}\right) \\ &\leq \exp((C + k)n) \left(1 + \alpha^{-\frac{n}{k-1}}\right) \end{aligned}$$

$$\leq \exp((C + k + 1)n)\alpha^{-\frac{n}{k-1}},$$

the last step being because $\alpha \leq 1$. Replacing $C + k + 1$ by C , we have deduced the upper bound.

For the lower bound, suppose now $\alpha \geq n/\binom{n}{k}$. Let $m = \lceil \frac{n}{\alpha} \rceil$ in the lower bound of Theorem 2.5. Since $\alpha \geq n/\binom{n}{k}$, we have $m \leq \binom{n}{k}$, so this is valid. We obtain that there are at least $\exp(C'n)\alpha^{-\frac{n}{k-1}}$ permutations in S_n containing at most $\lceil \frac{n}{\alpha} \rceil$ copies of π for some $C' = C'(\pi)$. Thus, by (1),

$$\begin{aligned} \mathbb{E}_\Lambda[|Av_{n,\Lambda}(\pi)|] &\geq \exp(C'n)\alpha^{-\frac{n}{k-1}}(1-\alpha)^{\lceil \frac{n}{\alpha} \rceil} \\ &\geq \exp(C'n)\alpha^{-\frac{n}{k-1}} \exp\left(-\frac{\alpha}{1-\alpha} \cdot \frac{2n}{\alpha}\right) \\ &\geq \exp\left(\left(C' - \frac{2}{1-\alpha}\right)n\right)\alpha^{-\frac{n}{k-1}}, \end{aligned}$$

where in the second line we used the inequality $\log(1-\alpha) \geq -\frac{\alpha}{1-\alpha}$ (an easy consequence of Taylor expansion) and $\lceil \frac{n}{\alpha} \rceil \leq \frac{2n}{\alpha}$ (immediate as $\alpha \leq 1$ and $n \geq 1$). This proves the lower bound with a constant of $C = C' - 4$ when $\alpha \leq \frac{1}{2}$.

With $\alpha \geq \frac{1}{2}$, note that $\mathbb{E}_\Lambda[|Av_{n,\Lambda}(\pi)|] \geq 1 \geq \exp(-n)\alpha^{-\frac{n}{k-1}}$ (as either the all-increasing or all-decreasing permutation avoids π over any hypergraph), so the constant -1 suffices. So letting $C = \min(C' - 4, -1)$ is sufficient to prove the lower bound, completing our argument. \square

4 Bounds on 0-1 Matrices

As in the proof strategy of [6], before we prove our result for permutations we first pass to the domain of 0-1 matrices. Since we would like to bound the number of permutations with few copies of π , we first show that a matrix M that contains few copies of the corresponding permutation matrix A_π must have few ones.

The technique we use to prove Theorem 2.3 is a classic method for proving supersaturation results; that is, we show that a random submatrix of M will with non-negligible probability contain at least one copy of A_π , so all of M must contain several copies of A_π .

In [6], Marcus and Tardos famously proved the following result, previously known as the Füredi-Hajnal conjecture.

Theorem 4.1 (Marcus-Tardos). *There exists a constant c_π such that for all n , any 0-1 matrix of size $n \times n$ containing at least $c_\pi n$ ones contains a copy of A_π .*

From this, we can immediately deduce the following (extremely weak) supersaturation result, which we will bootstrap using sampling into our stronger results.

Lemma 4.2. *With c_π as in Theorem 4.1, any 0-1 matrix of size $n \times n$ with m ones contains at least $m - c_\pi n$ copies of A_π .*

Proof of Lemma 4.2. We proceed by induction on m . For $m \geq c_\pi n$ the result is trivial.

Now take $m > c_\pi n$ and assume Lemma 4.2 for $m - 1$. Take a 0-1 matrix M of size $n \times n$ with m ones. Now, by Theorem 4.1, it contains a copy of A_π . Let M' be the matrix given by M with one of the ones in this copy of A_π changed to a 0. By the inductive hypothesis, M' contains at least $m - c_\pi n - 1$ copies of A_π . But in going from M to M' we eliminated at least one copy of A_π by design, so M must have at least $m - c_\pi n$ copies of A_π , finishing the induction. \square

We are now ready to prove Theorem 2.3.

Proof of Theorem 2.3. Take $k \in \mathbb{Z}^+$, $\pi \in S_k$, and let c_π be as given by Theorem 4.1 and Lemma 4.2. Let M be a 0-1 matrix of size $n \times n$ containing a ones, $Cn \leq a \leq n^2$, with C to be chosen later.

Take an r by r submatrix R of M , with r to be chosen later. Let the density of ones in R (that is, the number of ones in R divided by r^2) be $1(R)$. Similarly, let the density of A_π in R (that is, the number of copies of A_π in R divided by $\binom{r}{k}^2$) be $\pi(R)$. Define $1(M)$ and $\pi(M)$ similarly; in particular, $1(M) = \frac{a}{n^2} \geq \frac{C}{n}$ by assumption.

In this notation, Lemma 4.2 tells us that

$$\binom{r}{k}^2 \pi(R) \geq r^2 1(R) - c_\pi r,$$

or rearranging,

$$1(R) \leq \frac{\binom{r}{k}^2}{r^2} \pi(R) + \frac{c_\pi}{r}. \quad (2)$$

Now, let R be a *random* $r \times r$ submatrix of M (we choose a random subset of size r of the rows and similarly for the columns). Now, for each copy of A_π in M (defined by k rows and k columns), there is a $\frac{\binom{r}{k}^2}{\binom{n}{k}^2}$ probability that all rows and columns corresponding to this copy of A_π are chosen to be in R . Thus the expected number of copies of A_π in R is $\frac{\binom{r}{k}^2}{\binom{n}{k}^2}$ times the number of copies of A_π in M , and therefore

$$\mathbb{E}[\pi(R)] = \pi(M).$$

Similarly, each entry of M has equal probability of appearing in R , and so

$$\mathbb{E}[1(R)] = 1(M) \geq \frac{a}{n^2},$$

as by assumption M has at least a ones.

Now that we have $1(M)$ and $\pi(M)$ expressed in terms of $1(R)$ and $\pi(R)$, (2) applied to R will give an inequality between $1(M)$ and $\pi(M)$. Explicitly,

$$1(M) = \mathbb{E}[1(R)] \leq \mathbb{E} \left[\frac{\binom{r}{k}^2}{r^2} \pi(R) + \frac{c_\pi}{r} \right] = \frac{\binom{r}{k}^2}{r^2} \pi(M) + \frac{c_\pi}{r}. \quad (3)$$

We now have lower bounded $\pi(M)$ (which is a scaling of the number of copies of A_π in M) in terms of the number of ones in M . It only remains to optimize the value of our parameter r in order to obtain the desired lower bound.

Clearly, we need $1(M) > \frac{c_\pi}{r}$ for (3) to give any bound on $\pi(M)$ at all, so we choose $r = \left\lfloor \frac{3c_\pi}{1(M)} \right\rfloor$. We require $r \leq n$ (so that we can sample $r \times r$ submatrices), but this holds as long as $1(M) \geq \frac{3c_\pi}{n}$; that is, M has at least $3c_\pi n$ ones. Thus taking $C = 3c_\pi$ in the statement of Theorem 2.3 is sufficient to satisfy $r \leq n$. (This will be our only restriction on C .)

Note that since $c_\pi \geq 1$ and $1(M) \leq 1$, we have $r \geq \frac{2c_\pi}{1(M)}$. Thus $1(M) - \frac{c_\pi}{r} \geq \frac{1(M)}{2}$. Substituting into (3),

$$\begin{aligned} \frac{1(M)}{2} &\leq 1(M) - \frac{c_\pi}{r} \\ &\leq \frac{\binom{r}{k}^2}{r^2} \pi(M). \end{aligned}$$

Now, since $n \geq r$, and the function $\frac{\binom{a}{k}}{a^k}$ is increasing for $a > k$ (and 0 on integers less than k), we have that $\binom{r}{k} \leq \frac{r^k}{n^k} \binom{n}{k}$. Substituting this yields

$$\begin{aligned} 1(M) &\leq 2 \frac{\binom{r}{k}^2}{r^2} \pi(M) \\ &\leq 2 \frac{r^{2k-2}}{n^{2k}} \binom{n}{k}^2 \pi(M) \\ &\leq 2 \frac{(3c_\pi)^{2k-2}}{(1(M))^{2k-2} n^{2k}} \binom{n}{k}^2 \pi(M). \end{aligned}$$

Letting $C' = C'(\pi) := \frac{1}{2(3c_\pi)^{2k-2}}$, we have shown that

$$C' \cdot 1(M)^{2k-1} n^{2k} \leq \binom{n}{k}^2 \pi(M). \quad (4)$$

Now, $1(M) = \frac{a}{n^2}$ by definition. Furthermore, the right hand side of (4) is simply the number of copies of A_π in M (by the definition of $\pi(M)$). Thus we have shown that M contains at least $C' \frac{a^{2k-1}}{n^{2k-2}}$ copies of A_π , so taking this value of C' and $C = 3c_\pi$ (as above), we have proven our upper bound.

To show that this bound is sharp, take a and n with $Cn \leq a \leq n^2$. We may modify a and n by at most a constant factor so that $n|a$ and $a|n^2$. Now, suppose $\pi(1) > \pi(k)$ without loss of generality. Divide our $n \times n$ matrix M into blocks of side length $\frac{a}{n}$ (so there are $\frac{n^2}{a}$ blocks on each side). Consider the $\frac{n^2}{a}$ blocks along the main (upper left to lower right) diagonal. Fill each of these blocks with ones, and fill the rest of M with zeroes. The idea is to show that any copy of A_π must fall entirely into one of the blocks.

How many copies of A_π are contained in M ? Recall that a copy of A_π is given by a set of k entries of M which equal 1, say at indices $(i_1, j_1), \dots, (i_k, j_k)$, with $i_1 < \dots < i_k$ and the relative ordering of the j_k given by π .

There are a ones in M (as required) and so at most a choices for (i_1, j_1) . Let B be the $\frac{a}{n} \times \frac{a}{n}$ block containing (i_1, j_1) . Now, $i_1 < i_k$ and $j_1 > j_k$ (since $\pi(1) > \pi(k)$), so we are looking for a point to the lower-left of (i_1, j_1) . But since all blocks containing ones are on the main diagonal, (i_k, j_k) must be contained in B as well.

Now, for all r , we have that $i_1 \leq i_r \leq i_k$, and since B is the only block in its row containing ones, all other entries (i_r, j_r) must be contained in B . So for each of the remaining $k - 1$ entries (i_r, j_r) with $r > 1$, there are at most $|B| = \left(\frac{a}{n}\right)^2$ choices. So in total there are at most

$$a \left(\frac{a}{n}\right)^{2(k-1)} = \frac{a^{2k-1}}{n^{2k-2}}$$

copies of π in M . Since we only had to adjust a, n by a constant factor in the start, this proves the desired sharpness bounds, completing the proof of Theorem 2.3. \square

En route to the proof of Theorem 2.5 in the next section, we bound the number of 0-1 matrices containing few copies of A_π .

Proposition 4.3. *Let $k \in \mathbb{Z}^+$ with $k > 1$, $\pi \in S_k$ be fixed. There is a constant $C = C(\pi)$ such for all $m, n \geq 0$, the number of 0-1 matrices of size $n \times n$ containing at most m copies of A_π is at most*

$$\exp\left(C\left(n + \sqrt[2k-1]{mn^{2k-2}}\right)\right).$$

Remark 4.4. Notice that Proposition 4.3 immediately implies the same bound on the number of permutations of length n avoiding π (since each such permutation matrix yields a 0-1 matrix of size $n \times n$ avoiding A_π). Unfortunately, this bound is not as strong as the one we need to prove Theorem 2.5. However, we will be able to bootstrap Proposition 4.3 to prove the full result by a technique involving ‘contracting’ each permutation matrix containing few copies of π to a smaller 0-1 matrix, and then applying the Proposition to this smaller matrix.

The explanation for the expression $n + \sqrt[2k-1]{mn^{2k-2}}$ is that Theorem 2.3 implies that any 0-1 matrix M of size $n \times n$ containing at most m copies of A_π has at most $O\left(n + \sqrt[2k-1]{mn^{2k-2}}\right)$ ones. This is because if this does not hold, then letting a be the number of ones in M , we have that both $a \gg n$ and $\frac{a^{2k-1}}{n^{2k-2}} \gg m$, which contradicts the Theorem.

Proof. The proof here parallels Klazar’s proof that Füredi-Hajnal implies Stanley-Wilf given in [5]. The idea is to ‘contract’ any $n \times n$ matrix M containing few copies of A_π to an $n/2 \times n/2$ matrix M' by dividing M up into 2×2 boxes and assigning each box a 1 if and only if there are any ones in the box in M . If M contains few copies of A_π , M' must also, so by Theorem 2.3 it must contain few ones. But the number of possible values of M given M' is exponential in the number of ones of M' , so we may upper bound

the number of possible choices for M by the number of possible choices for M' . This will give a recursion yielding the desired bound.

In particular, let $S(n, m)$ be the set of 0-1 matrices of size $n \times n$ containing at most m copies of π , and let $f(n, m) = |S(n, m)|$. For a 0-1 matrix M of size $n \times n$ with $2|n$, let the 2-contraction of M be the $n/2 \times n/2$ 0-1 matrix M' such that $M'_{i,j} = 0$ if and only if $M_{2i-1,2j-1} = M_{2i-1,2j} = M_{2i,2j-1} = M_{2i,2j} = 0$.

Now, for each copy of A_π in M' , there is at least one corresponding copy of A_π in M . This is because a copy of A_π in M' corresponds to a choice of k 1-entries of M' with relative row- and column- ordering given by π , and each 1-entry of M' corresponds (in an order-preserving way) to at least one 1-entry of M . Thus M contains at least as many copies of A_π as its 2-contraction M' , so M' must also contain at most m copies of A_π .

Therefore, if $M \in S(n, m)$, then we must have $M' \in S(n/2, m)$, where M' is the 2-contraction of M . Thus

$$f(n, m) = |S(n, m)| \leq \sum_{M' \in S(n/2, m)} |\{M : M' \text{ is the 2-contraction of } M\}|. \quad (5)$$

Now, given a matrix M' , how many matrices M 2-contract to M' ? For every 0-entry of M' , the corresponding four entries of M must be 0, so there are no choices to be made. For every 1-entry of M' , the corresponding four entries of M may be either 1 or 0 (but not all 0), so there are 15 choices for those entries of M . Thus there are $15^{(\# \text{ of ones in } M')}$ matrices that 2-contract to M' . Combining this with (5), we obtain that

$$f(n, m) \leq \sum_{M' \in S(n/2, m)} 15^{(\# \text{ of ones in } M')} \leq f(n/2, m) \cdot \max_{M' \in S(n/2, m)} 15^{(\# \text{ of ones in } M')} \quad (6)$$

We now apply Theorem 2.3. For $M' \in S(n/2, m)$, we know that M' has at most m copies of A_π by definition, so by Theorem 2.3 it must have at most

$$O\left(n/2 + {}^{2k-1}\sqrt{m(n/2)^{2k-2}}\right) = O\left(n + {}^{2k-1}\sqrt{mn^{2k-2}}\right)$$

ones (by the discussion at the beginning of the proof). Substituting into (6),

$$f(n, m) \leq f(n/2, m) \cdot \exp(C_0(n + {}^{2k-1}\sqrt{mn^{2k-2}})).$$

for some $C_0 = C_0(\pi)$. This recursion is fairly easy to solve; we see that for $a \in \mathbb{Z}^{\geq 0}$

$$\begin{aligned} \log(f(2^a, m)) &\leq \log(f(1, m)) + C_0 \sum_{i=1}^a \left(2^i + {}^{2k-1}\sqrt{m2^{i(2k-2)}}\right) \\ &\leq 1 + C_0 \left(2^{a+1} + {}^{2k-1}\sqrt{m} \cdot \frac{2^{\frac{(a+1)(2k-2)}{2k-1}}}{2^{\frac{2k-2}{2k-1}} - 1}\right) \\ &\leq (C_0 + 1) \left(2^{a+1} + 2 \cdot {}^{2k-1}\sqrt{m} \cdot 2^{(a+1)(2k-2)}\right), \end{aligned}$$

where we simply summed the geometric series and used that $\log(f(1, m)) \leq \log(2) \leq 1$ and that $2^{\frac{2k-2}{2k-1}} \geq \frac{3}{2}$ for $k \geq 2$. Now, $f(n, m)$ is nondecreasing in n (as we may ‘pad’ any $n \times n$ matrix with zeroes to form an $n' \times n'$ matrix with the same number of copies of A_π , and this process is injective). For any n , take $a \in \mathbb{Z}^{\geq 0}$ such that $2^{a-1} < n \leq 2^a$. Then by the previous computation,

$$\begin{aligned} \log(f(n, m)) &\leq \log(f(2^a, m)) \\ &\leq (C_0 + 1) \left(2^{a+1} + 2 \cdot \sqrt[2k-1]{m \cdot 2^{(a+1)(2k-2)}} \right) \\ &\leq (C_0 + 1) \left(4n + 2 \cdot \sqrt[2k-1]{m \cdot (4n)^{2k-2}} \right) \\ &\leq 8(C_0 + 1) \left(n + \sqrt[2k-1]{mn^{2k-2}} \right). \end{aligned}$$

Letting $C = 8(C_0 + 1)$ completes the proof of Proposition 4.3. \square

Now that we have bounded the total number of 0-1 matrices that contain few copies of A_π , in the next section we may bound the number of permutations that contain few copies of π .

5 Permutations with Few Copies of π

This section will be devoted to the proof of Theorem 2.5.

5.1 Proof of the Upper Bound

Our proof of the upper bound of Theorem 2.5 will proceed in the following steps.

1. For a suitable $b \in \mathbb{Z}^+$ and for any matrix permutation σ containing few copies of π , take the b -contraction (analogous to the 2-contraction in the last section) of the matrix A_σ to get some $n/b \times n/b$ matrix B_σ .
2. B_σ is a 0-1 matrix containing few copies of A_π , so we can apply Proposition 4.3. Thus as σ ranges over all permutations in S_n containing few copies of π , B_σ ranges through only a small number of distinct matrices. Thus to bound the number of σ containing few copies of π , it suffices to bound for any matrix B the number of $\sigma \in S_n$ such that $B_\sigma = B$.
3. We bound the desired quantity $|\{\sigma \in S_n : B_\sigma = B\}|$ for any $n/b \times n/b$ matrix B using a simple counting argument.

The key is that since we are only applying Proposition 4.3 to an $n/b \times n/b$ matrix instead of an $n \times n$ matrix, we can obtain a much better bound (as long as b is chosen accordingly).

Step 1. Let

$$S_n(m, \pi) := \{\sigma \in S_n : \sigma \text{ contains at most } m \text{ copies of } \pi\}.$$

To prove the upper bound of Theorem 2.5, we would like to show that

$$|S_n(m, \pi)| \leq \exp(O(n)) \cdot \max \left(1, \left(\frac{m}{n} \right)^{\frac{n}{k-1}} \right).$$

First suppose $m < n$, such that the max is dominated by the first term. Then Proposition 4.3 guarantees that the number of 0-1 matrices of size $n \times n$ containing at most m copies of A_π is at most $\exp(O(n))$. Since each $\sigma \in S_n$ containing at most m copies of π gives rise to a permutation matrix A_σ that contains at most m copies of A_π , we see that the number of σ containing at most m copies of π is $\exp(O(n))$, as desired.

Now suppose $m \geq n$. Take $b = \sqrt[2k-2]{\frac{m}{n}}$. Just as we took the 2-contraction of a matrix in the proof of Proposition 4.3, we will define the b -contraction of any 0-1 matrix of size $n \times n$. The b -contraction of such a matrix A is the 0-1 matrix B such that the dimensions of B are $\lceil \frac{n}{b} \rceil \times \lceil \frac{n}{b} \rceil$, and such that $B_{i,j} = 1$ if and only if there exists i', j' with $\lceil \frac{i'}{b} \rceil = i$ and $\lceil \frac{j'}{b} \rceil = j$ such that $A_{i',j'} = 1$ (so if $A_{i',j'} = 0$ for all such i', j' , then $B_{i,j} = 0$). Let

$$n' := \lceil \frac{n}{b} \rceil = \lceil \sqrt[2k-2]{n^{2k-1}m^{-1}} \rceil$$

so that B is here an $n' \times n'$ matrix.

For all $\sigma \in S_n$, let B_σ be the b -reduction of A_σ .

Step 2. Similarly to the proof of Proposition 4.3 (with the 2-contraction), any occurrence of A_π in B_σ will correspond to at least one occurrence of A_π in A_σ . This again comes from, for each 1-entry in B appearing in that occurrence of A_π , choosing a corresponding 1-entry of A , and realizing that these 1-entries have the same relative row- and column-ordering.

Now, we have shown that each occurrence of A_π in B_σ gives rise to at least one occurrence of A_π in A_σ (it is easy to see that these occurrences are all distinct), and the occurrences of A_π in A_σ correspond to occurrences of π in σ . Thus for all $\sigma \in S_n(m, \pi)$, the matrix B_σ contains at most m copies of A_π .

By Proposition 4.3 (using the fact that $m > n \geq n'$), there are at most

$$\exp \left(C \left(\sqrt[2k-1]{mn'^{2k-2}} \right) \right)$$

matrices of dimension $n' \times n'$ that contain at most m copies of A_π (for $C = C(\pi)$). So as σ ranges over all elements of $S_n(m, \pi)$, B_σ ranges over at most

$$\begin{aligned} \exp \left(C \left(\sqrt[2k-1]{mn'^{2k-2}} \right) \right) &\leq \exp \left(C \left(\sqrt[2k-1]{m \left(2 \sqrt[2k-2]{n^{2k-1}m^{-1}} \right)^{2k-2}} \right) \right) \\ &= \exp \left(C \left(\sqrt[2k-1]{2^{2k-2}mn^{2k-1}m^{-1}} \right) \right) \\ &\leq \exp(2Cn) \end{aligned}$$

different matrices (where we used the fact that $n' = \lceil \sqrt[2k-2]{n^{2k-1}m^{-1}} \rceil \leq 2 \sqrt[2k-2]{n^{2k-1}m^{-1}}$ as $m \leq \binom{n}{k} < n^{2k-1}$). Therefore,

$$|S_n(m, \pi)| = \sum_{B \text{ size } n' \times n'} |\{\sigma \in S_n(m, \pi) : B_\sigma = B\}| \tag{7}$$

$$\leq \exp(2Cn) \cdot \max_{B \text{ size } n' \times n'} |\{\sigma \in S_n(m, \pi) : B_\sigma = B\}| \quad (8)$$

$$\leq \exp(2Cn) \cdot \max_{B \text{ size } n' \times n'} |\{\sigma \in S_n : B_\sigma = B\}|. \quad (9)$$

Step 3. It only remains to bound $\max_B |\{\sigma \in S_n : B_\sigma = B\}|$ from above for all 0-1 matrices B of size $n' \times n'$. That is, we must prove an upper bound on the number of permutation matrices of size $n \times n$ that b -contract to a particular matrix.

Now, since A_σ is a permutation matrix, it has n ones. By the definition of b -contraction, B_σ must have at most n ones. So in computing $\max_B |\{\sigma \in S_n : B_\sigma = B\}|$ we may assume B is an $n' \times n'$ matrix with at most n ones.

Let B be such a matrix, and suppose there are a_i ones in the i^{th} row of B . Then $\sum_{i=1}^{n'} a_i \leq n$. How many choices are there for σ such that $B_\sigma = B$? Consider the first row of A_σ , in which there is exactly one 1. This 1, when we take the b -reduction, must correspond to a 1 of B in the first row of B . There are a_1 such ones in the first row of B , and each one corresponds to at most $\lceil b \rceil$ entries in the first row of A_σ . Thus there are at most $\lceil b \rceil \cdot a_1$ ways to choose the position of the 1 in the first row of A_σ —in other words, to choose $\sigma(1)$.

Similarly, the 1-entry in the i^{th} row of A_σ must correspond to a 1-entry in the $\lceil \frac{i}{b} \rceil^{\text{th}}$ row of B , so there are at most $\lceil b \rceil \cdot a_{\lceil \frac{i}{b} \rceil}$ ways to choose the value of $\sigma(i)$. This implies that the total number of choices for σ such that $B_\sigma = B$ is at most

$$\prod_{i=1}^n \lceil b \rceil \cdot a_{\lceil \frac{i}{b} \rceil} = \lceil b \rceil^n \prod_{i=1}^n a_{\lceil \frac{i}{b} \rceil}. \quad (10)$$

Now, in the sum

$$\sum_{i=1}^n a_{\lceil \frac{i}{b} \rceil}, \quad (11)$$

every particular a_j occurs at most $\lceil b \rceil$ times, once for every i such that $bj - b < i \leq bj$.

Thus (11) is bounded by $\lceil b \rceil \sum_{j=1}^{n'} a_j \leq \lceil b \rceil \cdot n$. So by the AM-GM inequality,

$$\prod_{i=1}^n a_{\lceil \frac{i}{b} \rceil} \leq \lceil b \rceil^n.$$

Substituting into (10), we see that there are at most $\lceil b \rceil^{2n}$ choices for σ such that $B_\sigma = B$. Finally, substituting into (9), we have derived that

$$|S_n(m, \pi)| \leq \exp(2Cn) \lceil b \rceil^{2n}.$$

Now by definition, $b = \sqrt[k-2]{\frac{m}{n}}$, and $m \geq n$, so $b \geq 1$ and $\lceil b \rceil \leq 2b = 2 \sqrt[k-2]{\frac{m}{n}}$. Therefore,

$$\lceil b \rceil^{2n} \leq 4^n \left(\frac{m}{n}\right)^{\frac{n}{k-1}}.$$

This implies that

$$|S_n(m, \pi)| \leq \exp((2C + 2)n) \left(\frac{m}{n}\right)^{\frac{n}{k-1}},$$

and replacing $2C + 2$ by C finishes the proof of the upper bound in Theorem 2.5.

5.2 Proof of the Lower Bound

To prove the lower bound of Theorem 2.5, we must exhibit at least

$$\exp(-O(n)) \cdot \max\left(1, \left(\frac{m}{n}\right)^{\frac{n}{k-1}}\right)$$

permutations $\sigma \in S_n$ such that σ contains at most m copies of π .

Suppose without loss of generality that $\pi(1) > \pi(k)$. For $m \leq n$ the all-increasing permutation avoids π , so we get a lower bound of 1, which is sufficient.

Now suppose $m > n$. Note that $S_n(m, \pi)$ is nondecreasing in m and that changing m by at most a constant multiple does not change our desired lower bound by more than an exponential factor. Thus we may without loss of generality modify m by a constant multiple. In particular, we may assume without loss of generality that $\frac{m}{n}$ is a $(k-1)^{st}$ power, say $a^{k-1} = \frac{m}{n}$, $a \in \mathbb{Z}^+$.

Let $S_{n,a}$ be the set of permutations $\sigma \in S_n$ such that:

$\sigma(1), \dots, \sigma(a)$ is a permutation of $1, \dots, a$

$\sigma(a+1), \dots, \sigma(2a)$ is a permutation of $a+1, \dots, 2a$

\vdots

$\sigma\left(\left(\lfloor \frac{n}{a} \rfloor - 1\right)a + 1, \dots, \sigma\left(\lfloor \frac{n}{a} \rfloor a\right)\right)$ is a permutation of $\left(\lfloor \frac{n}{a} \rfloor - 1\right)a + 1, \dots, \lfloor \frac{n}{a} \rfloor a$

$\sigma\left(\lfloor \frac{n}{a} \rfloor a + 1, \dots, n\right)$ is a permutation of $\lfloor \frac{n}{a} \rfloor a + 1, \dots, n$.

It suffices to prove that

1. $|S_{n,a}| \geq \exp(-O(n)) \left(\frac{m}{n}\right)^{\frac{n}{k-1}}$, and
2. $|S_{n,a}| \in S_n(m, \pi)$; that is, any element of $S_{n,a}$ contains at most m copies of π .

Let $n = qa + r$, $q, r \in \mathbb{Z}^{\geq 0}$, $r < a$. Then $|S_{n,a}| = (a!)^q \cdot r!$. Since $t! \geq \left(\frac{t}{e}\right)^t$ for all $t \in \mathbb{Z}^{\geq 0}$ (using $0! = 1$), we see that

$$\begin{aligned} |S_{n,a}| &\geq \left(\frac{a}{e}\right)^{qa} \left(\frac{r}{e}\right)^r \\ &= \left(\frac{r}{a}\right)^r \left(\frac{a}{e}\right)^n, \end{aligned}$$

as $qa + r = n$. Now, the function x^x is minimized for $x \in [0, 1]$ when $x = \frac{1}{e}$, so $x^x \geq e^{-\frac{1}{e}}$. Thus $\left(\frac{r}{a}\right)^r = \left(\frac{r}{a}\right)^{\frac{r}{a}a} \geq \exp\left(-\frac{r}{e}\right)$. Now, $m \leq \binom{n}{k} < n^k$, and therefore $a < n$. Thus

$$\left(\frac{r}{a}\right)^r \geq \exp(-n).$$

Therefore,

$$|S_{n,a}| \geq a^n \exp(-2n) = \exp(-2n) \left(\frac{m}{n}\right)^{\frac{n}{k-1}}.$$

This gives is our desired bound on $|S_{n,a}|$. It thus suffices to show that any element of $S_{n,a}$ contains at most m copies of π .

Suppose $\sigma \in S_{n,a}$. At what indices can π occur in σ ? Let π occur at some set of k indices $i_1 < \dots < i_k$. Then since $\pi(1) > \pi(k)$, we must have $\sigma(i_1) > \sigma(i_k)$, while of course $i_1 < i_k$. By the definition of $S_{n,a}$, this can only occur when $\lceil \frac{i_1}{a} \rceil = \lceil \frac{i_k}{a} \rceil$. Since $i_1 < \dots < i_k$, this means that there is some t , $0 \leq t \leq q$, such that $ta + 1 \leq i_1 < \dots < i_k \leq (t+1)a$ (where again $qa + r = n$, $r < a$).

Given a particular value of t , there are thus at most $\binom{a}{k}$ choices for (i_1, \dots, i_k) . However, if $t = q$, we have that $qa + 1 \leq i_1 < \dots < i_k \leq qa + r = n$, so there are in this case only at most $\binom{r}{k}$ choices for (i_1, \dots, i_k) . Thus the total number of occurrences of π in σ is at most

$$\begin{aligned} q \binom{a}{k} + \binom{r}{k} &< qa^k + r^k \\ &\leq qa^k + ra^{k-1} \\ &\leq (qa + r)a^{k-1} \\ &= na^{k-1} \\ &= m. \end{aligned}$$

Thus $S_{n,a} \subseteq S_n(m, \pi)$, so we have proved the lower bound and we are done.

6 The Fixed Hypergraph Case

For fixed positive integer k , let Λ be a k -uniform hypergraph on n vertices satisfying the preconditions of Theorem 2.8. That is, for some L , Λ contains a collection of L -vertex cliques where each of the n vertices belongs to at least $\delta(\Lambda) \geq 1$ cliques in the collection and at most $\Delta(\Lambda) = O(1)$.

We would like to show that for every permutation $\pi \in S_k$,

$$Av_{n,\Lambda}(\pi) = O\left(\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n\right)$$

for all $\epsilon > 0$.

For $L = \Theta(n^c)$ with $c \in (0, 1]$, this bound is a strict improvement on the $n!$ total n -permutations. For $L = n$, we match the Stanley-Wilf conjecture up to a log-exponential factor $\log(n)^{O(n)}$, which is asymptotically dominated by the linear exponential term $n^{O(n)}$ of the conjecture.

It may seem unnatural at first to restrict our arguments only to hypergraphs containing polynomially large cliques. However, we see that there are very dense hypergraphs Λ^* with $O(1)$ maximal clique size for which the number of n -permutations π that are Λ^* -avoiding is $O(n)^n$. For example, consider partitioning the vertices of Λ^* into two parts,

$\{1, \dots, n/2\}$ and $\{n/2 + 1, \dots, n\}$, and adding an edge to Λ^* for every collection of k vertices not entirely lying in a single part. This graph will be very dense, containing $\binom{n}{k} - 2\binom{n/2}{k} \approx (1 - \frac{1}{2^k})\binom{n}{k}$ edges. However, there is a large class of n -permutations avoiding π on these edges. Say, without loss of generality, that $\pi(1) < \pi(k)$. We see that all n -permutations σ in which the $n/2$ largest elements belong in the first $n/2$ indices and the $n/2$ smallest elements belong in the last $n/2$ indices necessarily Λ^* -avoid π . Each edge of Λ^* corresponds to a sub permutation $(\sigma(x_1), \dots, \sigma(x_k))$ in which $\sigma(x_1) > \sigma(x_k)$ and so it cannot be a copy of π . There are $(n/2)!^2 \approx (\frac{n}{2e})^n = O(n)^n$ such permutations, and so there is no meaningful bound we can prove on the number of Λ^* -avoidant n -permutations.

Importantly, multipartite graphs are characterized by their small maximal cliques. The bipartite graph we considered has maximal clique $2(k - 1)$, taking $k - 1$ vertices from each part. Thus, our bounds on Λ -avoidance being contingent on Λ containing large cliques is necessary. Though, the preconditions involving the minimal $\delta(\Lambda)$ and maximal $\Delta(\Lambda)$ number of cliques per vertex is more of an artifact of the proof technique. A more general argument with a slightly less specific precondition is likely to hold.

7 Hypergraph Formulation of Pattern-Avoidance

We consider a k -uniform hypergraph H on an $n \times n$ grid of vertices $V(H)$, which we index $v(i, j)$. Define a canonical set to be a subset of $V(H)$ of size n containing exactly one vertex from each row and each column. We see that a canonical set corresponds bijectively to an n -permutation σ . For a k -permutation π , we add edges to H in such a way that each canonical set is independent if and only if its corresponding n -permutation is Λ -avoidant of π . Essentially, we add an edge for each copy of π in the vertices on columns in $E(\Lambda)$. For all $1 \leq x_1 < x_2 < \dots < x_k \leq n$ with $\{x_1, \dots, x_k\} \in E(\Lambda)$ and all $1 \leq y_1 < \dots < y_k \leq n$, we have $\{v(x_1, y_{\pi(1)}), v(x_2, y_{\pi(2)}), \dots, v(x_k, y_{\pi(k)})\} \in E(H)$. We see that a canonical set containing the vertices of this edge would correspond to a permutation σ that contains a copy of π at indices x_1, \dots, x_k , as desired.

We want to show that the number of n -permutations that Λ -avoid π is $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$. Since each permutation corresponds to a single canonical set, we want to show that the number of independent canonical sets is $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$. In fact, our goal will be to prove a stronger claim, that the number of independent sets of size n , of which the independent canonical sets are a subset, is $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$.

8 The Hypergraph Containers Lemma

We introduce a version of the hypergraph container lemma due to Balogh, Morris, and Samotij [1]. Essentially, the container lemma is a means of placing the vertices of a hypergraph into a collection of containers \mathcal{C} in such a way that each independent set in the hypergraph belongs to one of the containers. Additionally, we ensure that no individual container contains too many vertices and that the number of containers isn't

too large. We let $\Delta_\ell(\mathcal{H})$ be the maximum number of hyperedges of \mathcal{H} that contain a given set of ℓ vertices.

Proposition 8.1 ([1] Theorem 2.2). *Let \mathcal{H} be a k -uniform hypergraph and let K be a constant. There exists a constant $g \in (0, 1)$ depending only on k and K such that the following holds. Suppose that for some $p \in (0, 1)$ and all $\ell \in \{1, \dots, k\}$,*

$$\Delta_\ell(\mathcal{H}) \leq K \cdot p^{\ell-1} \cdot \frac{e(\mathcal{H})}{v(\mathcal{H})}$$

Then, there exists a family $\mathcal{C} \subseteq \mathcal{P}(V(\mathcal{H}))$ of containers with the following properties:

1. $|\mathcal{C}| \leq \binom{v(\mathcal{H})}{\leq kp v(\mathcal{H})} \leq \left(\frac{e}{kp}\right)^{kp v(\mathcal{H})}$,
2. $|G| \leq (1 - g) \cdot v(\mathcal{H})$ for each $G \in \mathcal{C}$,
3. each independent set of \mathcal{H} is contained in some $G \in \mathcal{C}$.

This lemma is extremely useful in bounding the number of independent sets of a hypergraph, as the number of independent sets is upper bounded by the sum of the number of independent sets in each container. Or, in our context, the number of independent sets of size n in H is upper bounded by the total number of independent sets of size n over all the containers. However, a single application of the container lemma to our problem will not be strong enough for our purposes, as a single container can still contain $(1 - g)|V(H)| = (1 - g)n^2$ vertices and potentially have $\binom{(1-g)n^2}{n} = O(n)^n$ many independent sets of size n . So, we will apply the lemma recursively. Each time we encounter a container with too many vertices, we apply the lemma to the subgraph induced by the vertices of the container and further break it up into more containers. We do this until all the containers are sufficiently small. Namely, we will attempt to apply the container lemma recursively until all the containers have $\leq U = \frac{Cn^2 \log^{2+\epsilon} n}{L}$ vertices, for a constant C that will only be in terms of k and π . Once the containers are this small, a naive upper bound will give us that the number of size- n independent sets in a container is at most $\binom{U}{n} = O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$.

Unfortunately, since we know nothing about the structure of the containers, we have no guarantee that the necessary Δ_ℓ bounds will hold, which are required to apply the lemma to a container. To overcome this problem we employ a strategy similar to that used by Morris and Saxton [7]. Consider a subgraph G of H induced by some subset/container of the vertices. If we remove some of the edges of G to produce a new subgraph G' , then every independent set of G will also be an independent set of G' . So, if we apply the containers lemma to G' , the resulting containers will also cover all the independent sets of G . This will be our approach: for each container-induced subgraph G with more than U vertices, we construct a subgraph $G' \subseteq G$ on the same vertex set with some edges of G removed. We construct G' to satisfy the preconditions of Proposition 8.1 for a sufficiently small p , and so, we can break up G' into containers using the proposition and recurse. In this recursive process, we guarantee that all of the independent sets in the original H are

preserved. We also ensure that we will not have too many containers in the end because we keep p small.

Formally, in Section 9, we prove the following lemma

Lemma 8.2. *Let $\gamma = \frac{1}{1-g}$, where g is defined in Proposition 8.1. For the hypergraph H defined in Section 7, consider a subgraph $G \subseteq H$ induced by some subset of the vertices, where*

$$Cn\gamma^{t-1} < |V(G)| \leq Cn\gamma^t$$

for some constant C and some $t \geq t_0 + 1$ with $Cn\gamma^{t_0} = U$. There exists a subgraph $G' \subseteq G$ on the same vertex set such that

$$\Delta_\ell(G') \leq K \cdot p_t^{\ell-1} \cdot \frac{|E(G')|}{|V(G)|}$$

for all $\ell \in \{1, \dots, k\}$, where $p_t = \frac{n}{t^{2+\epsilon}|V(G)|}$

This lemma enables the proof of Theorem 2.8 using the recursive hypergraph containers strategy.

Theorem 2.8 For every permutation π and any $\epsilon > 0$, the number of n -permutations Λ -avoiding π is $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$ as long as Λ is k -uniform and satisfies the following:

Λ contains a collection of L -vertex cliques where each of the n vertices belongs to at least $\delta(\Lambda) \geq 1$ cliques in the collection and at most $\Delta(\Lambda) = O(1)$.

Proof of Theorem 2.8. As stated earlier in this section, to prove Theorem 2.8 it is sufficient to prove that the hypergraph H has at most $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$ independent sets of size n .

Lemma 8.2 shows that, for a general container G , we can apply the container lemma for a certain $p = p_t$ depending on the size of G , and further split G into more containers. Starting from the original graph H , we can repeat this process recursively until all of our containers have $\leq U = \frac{Cn^2 \log^{2+\epsilon} n}{L}$ vertices. We are trying to count the number of independent sets of size n in the original hypergraph and we know every independent set in the original graph is a subset of one of these containers. Each container of size $\leq U$ has $\leq \binom{U}{n} \leq \left(\frac{eU}{n}\right)^n = O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$ subsets of size n , and so the number of independent sets of size n in this container is also bounded by this amount. Therefore, all that remains to show is that the number of containers is singly exponential in n . That is, upper bounded by c^n for some $c = O(1)$. If we can show this, then we will have

$$\begin{aligned} & \text{(number of size-}n\text{ independent sets in } H) \\ & \leq \sum_{\text{containers } C} \text{(number of size-}n\text{ independent sets in } C) \\ & = \text{(number of containers)} \cdot O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n \end{aligned}$$

$$= c^n \cdot O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n = O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$$

Say somewhere along this branching recursive process, we encounter a container G that we want to split into further containers with $Cn\gamma^{t-1} < v(G) \leq Cn\gamma^t$ and $t \geq t_0 + 1$. From Lemma 8.2, we know we can apply the container lemma with $p = p_t = \frac{n}{t^{2+\epsilon}|V(G)|}$ and split G into at most

$$\left(\frac{e}{kp_t}\right)^{kp_tv(G)} = \left(\frac{et^{2+\epsilon}|V(G)|}{kn}\right)^{kn/t^{2+\epsilon}} \leq \left(\frac{et^{2+\epsilon}Cn\gamma^t}{kn}\right)^{kn/t^{2+\epsilon}} = \left(\frac{et^{2+\epsilon}C\gamma^t}{k}\right)^{kn/t^{2+\epsilon}}$$

containers. Additionally, we know that all of the resulting containers will contain at most $(1 - \gamma)v(G) \leq Cn\gamma^{t-1}$ vertices. We will subsequently break down these child containers using $p = p_s$ for some $s \leq t - 1$. Say $T = \log_\gamma(n/C)$ or equivalently $Cn\gamma^T = n^2$. In the worst case, after we break up H with $p = p_T$, we break up all of H 's child containers with $p = p_{T-1}$, all of H 's grandchild containers with $p = p_{T-2}$, etc all the way to $p = p_{t_0+1}$. However, we can never encounter two consecutive generations of containers on which we apply the containers lemma with the same p_t ; t is always strictly decreasing. Thus, the number of containers we have at the end is at most

$$\prod_{t=t_0+1}^T \left(\frac{e}{kp_t}\right)^{kp_tv(G)} \leq \prod_{t=t_0+1}^T \left(\frac{et^{2+\epsilon}C\gamma^t}{k}\right)^{kn/t^{2+\epsilon}}$$

Defining $A = \gamma \left(\frac{et^{2+\epsilon}C}{k}\right)^{1/t}$,

$$= \prod_{t=t_0+1}^T (A^t)^{kn/t^{2+\epsilon}} = \left(A^{k \sum_{t=t_0+1}^T \frac{1}{t^{1+\epsilon}}}\right)^n$$

We note that $A^{k \sum_{t=t_0+1}^T \frac{1}{t^{1+\epsilon}}} = O(1)$ as $\sum_{t=1}^\infty \frac{1}{t^{1+\epsilon}}$ is a convergent sum and $t^{(2+\epsilon)/t}$ has a finite upper bound for all t . Moreover, A is a constant in terms of C and γ which are only in terms of π and k , and have no dependency on n . Thus, the number of containers at the end of the branching process is singly exponential in n , as desired. \square

9 Block Decomposition of G

We now demonstrate how to construct a subgraph G' of G that satisfies the preconditions of the containers lemma and enables the recursive argument.

Lemma. 8.2 Let $\gamma = \frac{1}{1-g}$, where g is defined in Proposition 8.1. For the hypergraph H defined in Section 7, consider a subgraph $G \subseteq H$ induced by some subset of the vertices, where

$$Cn\gamma^{t-1} < |V(G)| \leq Cn\gamma^t$$

for some constant C and some $t \geq t_0 + 1$ with $Cn\gamma^{t_0} = U$. There exists a subgraph $G' \subseteq G$ on the same vertex set such that

$$\Delta_\ell(G') \leq K \cdot p_t^{\ell-1} \cdot \frac{|E(G')|}{|V(G)|}$$

for all $\ell \in \{1, \dots, k\}$, where $p_t = \frac{n}{t^{2+\epsilon}|V(G)|}$

Proof. Note that, for any G' we construct, we will have $\Delta_k(G') = 1$ as a set of k vertices belonging to multiple edges would imply that we have duplicate edges. So, in order to satisfy the condition for $\ell = k$, we have $1 \leq K \cdot p_t^{k-1} \cdot \frac{|E(G')|}{|V(G)|}$, and so we must have $|E(G')| \geq \frac{|V(G)|}{Kp_t^{k-1}}$. We define

$$N = \frac{|V(G)|}{Kp_t^{k-1}}$$

and so, we will construct a G' with $|E(G')| \geq N$. Therefore, it will be sufficient to construct a G' satisfying

$$\Delta_\ell(G') \leq K \cdot p_t^{\ell-1} \cdot \frac{N}{|V(G)|} = \frac{1}{p_t^{k-\ell}} \tag{12}$$

for all $1 \leq \ell \leq k$.

To construct a subset of the edges of G satisfying (12), we must take advantage of some structural regularity of G . However, on the whole, all we know about G is roughly how many vertices it contains. Yet, we can take advantage of the fact that, specifically on the vertices of G corresponding to some clique in Λ , the edges are well-behaved and will lend themselves to an intricate construction. Namely, for each of the L -cliques in Λ , we define its ‘‘block’’ to be the subgraph of G induced by the set of vertices in $V(G)$ belonging to the L columns corresponding to this clique. We will construct G' by constructing a block subgraph B' for every block B and then taking G' to be the union of all the B' .

We call a block B rich if

$$|V(B)| \geq d = \sqrt{Lt^{2+\epsilon}|V(G)|}$$

We will show that for every rich block B , there exists a subgraph B' of B , on the same vertex set with

$$|E(B')| = N_B = \frac{2\Delta(\Lambda)}{\delta(\Lambda)} \cdot \frac{|V(B)|}{Kp_t^{k-1}}$$

We will also show that B' can be chosen so that

$$\Delta_\ell(B') \leq \frac{1}{\Delta(\Lambda)p^{k-\ell}}$$

If we can prove that such a B' exists for every rich block B , then we can construct G' by taking the union of all the B' . We see that, for any collection of ℓ vertices v_1, \dots, v_ℓ ,

$$\deg_{G'}(v_1, \dots, v_\ell) \leq \sum_{\text{rich blocks } B} \deg_{B'}(v_1, \dots, v_\ell) \leq \frac{\Delta(\Lambda)}{\Delta(\Lambda)p^{k-\ell}}$$

since any collection of ℓ vertices, as well as any single vertex, belongs to at most $\Delta(\Lambda)$ blocks. And so, $\Delta_\ell(G') \leq \frac{1}{p_t^{k-\ell}}$ as desired. We also see that we will have at least N edges in the union of the B' because

$$\begin{aligned} \left| \bigcup_{\text{rich blocks } B} E(B') \right| &\geq \frac{1}{\Delta(\Lambda)} \sum_{\text{rich blocks } B} |E(B')| \\ &= \frac{1}{\Delta(\Lambda)} \sum_{\text{rich blocks } B} N_B \\ &= \frac{1}{\Delta(\Lambda)} \sum_{\text{rich blocks } B} \left(\frac{2\Delta(\Lambda)}{\delta(\Lambda)Kp_t^{k-1}} \right) |V(B)| \\ &= \frac{2}{\delta(\Lambda)Kp_t^{k-1}} \cdot \sum_{\text{rich blocks } B} |V(B)| \end{aligned}$$

and we see

$$\sum_{\text{rich blocks } B} |V(B)| = \sum_{\text{blocks } B} |V(B)| - \sum_{\text{unrich blocks } B} |V(B)|$$

where

$$\sum_{\text{blocks } B} |V(B)| \geq \delta(\Lambda)|V(G)|$$

since each vertex belongs to at least $\delta(\Lambda)$ blocks, and

$$\sum_{\text{unrich blocks } B} |V(B)| \leq d(\text{number of unrich blocks}) \leq d(\text{number of blocks})$$

Now, since $Cn\gamma^{t-1} < |V(G)| \leq n^2$

$$d = \sqrt{Lt^{2+\epsilon}|V(G)|} < \sqrt{L|V(G)| \log_\gamma^{2+\epsilon} \left(\frac{\gamma|V(G)|}{Cn} \right)} \leq \sqrt{L|V(G)| \log_\gamma^{2+\epsilon} n}$$

for $C \geq \gamma$. And since each of the n vertices in Λ belongs to at most $\Delta(\Lambda)$ of the size L cliques, the number of L -cliques, which is the number of blocks, is at most $\Delta(\Lambda)n/L$. So,

$$\sum_{\text{rich blocks } B} |V(B)| \geq \sum_{\text{blocks } B} |V(B)| - d(\text{number of blocks})$$

$$\begin{aligned} &\geq \delta(\Lambda)|V(G)| - \left(\sqrt{L|V(G)| \log_\gamma^{2+\epsilon} n} \right) (\Delta(\Lambda)n/L) \\ &\geq \delta(\Lambda)|V(G)|/2 \end{aligned}$$

because

$$\begin{aligned} |V(G)| &\geq U = \frac{Cn^2 \log^{2+\epsilon} n}{L} \\ \therefore \sqrt{|V(G)|} &\geq \sqrt{\frac{Cn^2 \log^{2+\epsilon} n}{L}} \\ \therefore |V(G)| &\geq \frac{n}{L} \cdot \sqrt{|V(G)| \cdot CL \log^{2+\epsilon} n} \\ \therefore \delta(\Lambda)|V(G)|/2 &\geq \left(\sqrt{L|V(G)| \log_\gamma^{2+\epsilon} n} \right) (\Delta(\Lambda)n/L) \end{aligned}$$

for $C \geq \frac{\Delta(\Lambda)^2}{(\delta(\Lambda)/2)^2 \log(\gamma)}$, which is not in terms of n and is therefore a valid bound on the constant C . And so,

$$\begin{aligned} \left| \bigcup_{\text{rich blocks } B} E(B') \right| &\geq \frac{2}{\delta(\Lambda)Kp_t^{k-1}} \cdot \sum_{\text{rich blocks } B} |V(B)| \\ &\geq \frac{2}{\delta(\Lambda)Kp_t^{k-1}} \cdot \delta(\Lambda)|V(G)|/2 \\ &= N \end{aligned}$$

as desired. □

10 Supersaturation on the Rich Blocks

From the previous section, we showed that, to prove Lemma 8.2, it was sufficient to show the following lemma about rich blocks. Again, we define blocks to be the subgraph induced by the vertices of G belonging to a certain collection of L columns of the $n \times n$ grid. These L columns represent a clique in the avoidance hypergraph Λ . So, for any k of these L columns $x_1 < x_2 < \dots < x_k$ and any $1 \leq y_1 < \dots < y_k \leq n$, we have $\{v(x_1, y_{\pi(1)}), v(x_2, y_{\pi(2)}), \dots, v(x_k, y_{\pi(k)})\} \in E(H)$. We will have an edge in our block B for every collection of k such vertices that belong to G . Now, we state the lemma.

Lemma 10.1. *For a block subgraph $B \subseteq G \subseteq H$ with*

$$|V(B)| \geq d = \sqrt{Lt^{2+\epsilon}|V(G)|}$$

and

$$Cn\gamma^{t-1} < |V(G)| \leq Cn\gamma^t$$

for some $t \geq t_0 + 1$, there exists a subgraph $B' \subseteq B$ on the same vertex set such that

$$|E(B')| = N_B = \frac{2\Delta(\Lambda)}{\delta(\Lambda)} \cdot \frac{|V(B)|}{Kp_t^{k-1}} \quad (13)$$

and

$$\Delta_\ell(B') \leq \frac{1}{\Delta(\Lambda)p^{k-\ell}} \quad (14)$$

for all $\ell \in \{1, \dots, k\}$, where $p_t = \frac{n}{t^{2+\epsilon}|V(G)|}$ and $\gamma = \frac{1}{1-g}$, where g is defined in Proposition 8.1.

Proof. From our definition of N_B in (13), we can rewrite (14) as

$$\Delta_\ell(B') \leq \frac{\delta(\Lambda)}{2\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1}N_B}{|V(B)|}$$

We start our construction of B' with the hypergraph B_0 on the vertices of B with no edges. We then iteratively construct B_1, B_2, \dots, B_{N_B} where we construct B_{i+1} by adding an edge to B_i . B_{N_B} will be our B' .

For every $\ell \in [1, k-1]$ and every $i \in [0, N_B-1]$, we define the dangerous set $D_\ell(B_i)$ to be the set of all sets of ℓ vertices $\{v_1, \dots, v_\ell\}$ where

$$|\{E \in E(B_i) \mid \{v_1, \dots, v_\ell\} \subseteq E\}| \geq \frac{\delta(\Lambda)}{2\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1}N_B}{|V(B)|} - 1$$

Now, we say that an edge $E \in E(B)$ is i -safe if $F \notin D_{|F|}(B_i)$ for every nonempty, strict subset $F \subset E$. Our goal for all i will be to construct B_{i+1} by adding an i -safe edge to B_i that is not already in $E(B_i)$. If this is always possible, we see that, for all $\ell \in \{1, \dots, k-1\}$,

$$\Delta_\ell(B_{i+1}) \leq \max \left(\Delta_\ell(B_i), \frac{\delta(\Lambda)}{2\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1}N_B}{|V(B)|} - 1 + 1 \right)$$

and therefore, we can show inductively that the B_{N_B} we construct will satisfy $\Delta_\ell(B_{N_B}) \leq \frac{\delta(\Lambda)}{2\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1}N_B}{|V(B)|}$ and be a valid choice for B' , as desired.

In order to show there is always an i -safe edge E not already in $E(B_i)$, it is sufficient to show that the number of i -safe edges is $\geq N_B$, meaning that, by the pigeonhole principle, one of them is not already in $E(B_i)$. Let Z be the number of i -safe edges in B . We want to show $Z \geq N_B$. The vertices of B belong to an $n \times L$ matrix grid. We define S to be the set of vertices in B that belong to a random submatrix, selecting each column independently with probability q and each row independently with probability $\frac{Lq}{n}$, for a fixed $q \in (0, 1]$. In expectation, we select qL rows and qL columns. Also note that the probability that a single vertex is included in S is $\frac{Lq^2}{n}$ as both its row and column need to be selected.

Then, we generate another vertex subset $S' \subseteq S$. We start with $S' = S$ and iteratively scan S' for subsets of vertices $F \in D_{|F|}(B_i)$. If we find such an F , we remove one of the vertices in F from S' and restart the scanning process. We terminate once no subset of S' belongs to $D_{|F|}(B_i)$. We note that the total number of vertices of S that we deleted is at most the number of $F \in D_{|F|}(B_i)$ contained in S originally, as some vertex deletions may have destroyed multiple dangerous F .

Now, we consider the subgraph R induced by S' and define the random variable X to be the number of i -safe edges in R . Since we removed a vertex from every dangerous F in S' , there will be no dangerous F in $V(R)$ and every edge in R is i -safe. So, we have $X = |E(R)|$.

The probability that any i -safe edge in B belongs to R is $\leq \left(\frac{Lq^2}{n}\right)^k$, as each of the k vertices in the edge belongs to S with probability $\frac{Lq^2}{n}$ independently, as they all occupy separate rows and columns. So, by linearity of expectation, we can upper bound

$$\mathbb{E}[X] \leq Z \left(\frac{Lq^2}{n}\right)^k$$

Now, to show that $Z \geq N_B$, it remains to give a sufficiently strong lower bound for $\mathbb{E}[X]$. The main tool for doing so is the Füredi-Hajnal conjecture (proved in 2004) which also formed the backbone of the original proof of the Stanley-Wilf conjecture. The Füredi-Hajnal conjecture [4] states that any 0-1 matrix A of size $x \times x$ that avoids a permutation matrix P can have at most $c_P x$ 1-entries, for a constant c_P only in terms of P . The direct implication of this theorem in the hypergraph setting is this. For a hypergraph with an $x \times x$ grid of vertices, and edges corresponding to the copies of P on this grid, any independent set of this graph has at most $c_P x$ vertices. Using this, we can lower bound the number of edges in R using a supersaturation argument.

Let $x = \max(\text{number of rows selected in } S, \text{number of columns selected in } S)$. So, all of the vertices in R belong to an $x \times x$ subgrid. We claim that, by Füredi-Hajnal, $|E(R)| \geq |V(R)| - c_P x$. While R has more than $c_P x$ vertices, we can find an edge in R and delete one of the vertices in that edge. This decreases $|V(R)|$ by 1, and decreases $|E(R)|$ by at least 1. Repeating this process until the number of vertices left in R is $c_P x$, we must have removed at least $|V(R)| - c_P x$ edges which were originally in R . Thus, by linearity of expectation,

$$\mathbb{E}[|E(R)|] \geq \mathbb{E}[|V(R)| - c_P x] = \mathbb{E}[|V(R)|] - c_P \mathbb{E}[x]$$

Now,

$$\begin{aligned} \mathbb{E}[|V(R)|] &= \mathbb{E}[|S'|] = \mathbb{E}[|S| - \text{at most 1 for each dangerous set in } S] \\ &\geq \mathbb{E}[|S|] - \sum_{\ell=1}^{k-1} \sum_{F \in D_\ell(B_i)} \Pr[F \subseteq S] \\ &= \frac{Lq^2}{n} |V(B)| - \sum_{\ell=1}^{k-1} |D_\ell(B_i)| \cdot \left(\frac{Lq^2}{n}\right)^\ell \end{aligned}$$

and

$$\begin{aligned}\mathbb{E}[x] &= \mathbb{E}[\max(\text{number of rows selected}, \text{number of columns selected})] \\ &\leq \mathbb{E}[\text{number of rows selected}] + \mathbb{E}[\text{number of columns selected}] \\ &= \frac{Lq}{n} \cdot n + q \cdot L = 2qL\end{aligned}$$

Therefore,

$$Z \left(\frac{Lq^2}{n} \right)^k \geq \mathbb{E}[|E(R)|] \geq \frac{Lq^2}{n} |V(B)| - \sum_{\ell=1}^{k-1} |D_\ell(B_i)| \cdot \left(\frac{Lq^2}{n} \right)^\ell - 2qc_P L$$

We have that $|V(B)| > U = Cn\gamma^{t_0}$ for some constant C . So, we take $C > 4c_P$, which is only in terms of π and is therefore a valid constraint on C . Setting $q = \frac{4c_P n}{|V(B)|} < 1$, we have

$$\frac{Lq^2}{n} |V(B)| - 2qc_P L \geq \frac{Lq^2}{2n} |V(B)|$$

and

$$Z \left(\frac{Lq^2}{n} \right)^k \geq \mathbb{E}[|E(R)|] \geq \frac{Lq^2}{2n} |V(B)| - \sum_{\ell=1}^{k-1} |D_\ell(B_i)| \cdot \left(\frac{Lq^2}{n} \right)^\ell$$

So, in order to show $Z \geq N_B$, it is sufficient to show

$$\frac{Lq^2}{2n} |V(B)| - \sum_{\ell=1}^{k-1} |D_\ell(B_i)| \cdot \left(\frac{Lq^2}{n} \right)^\ell \geq N_B \left(\frac{Lq^2}{n} \right)^k = \frac{2\Delta(\Lambda)}{\delta(\Lambda)} \cdot \frac{|V(B)|}{Kp_t^{k-1}} \left(\frac{Lq^2}{n} \right)^k \quad (15)$$

We can bound $|D_\ell(B_i)|$ by double counting F, E pairs where

$$F = \{v_1, \dots, v_\ell\} \subseteq E \in E(B_i)$$

For an upper bound, we know there are $i \leq N_B$ ways to choose $E \in E(B_i)$ and there are $\binom{k}{\ell} \leq 2^k$ ways to choose an F belonging to that E . For a lower bound, each $F \in D_\ell(B_i)$ belongs to at least $\frac{\delta(\Lambda)}{2\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1} N_B}{|V(B)|} - 1$ many edges and each $F \notin D_\ell(B_i)$ belongs to at least 0 edges. So,

$$\begin{aligned}2^k N_B &\geq \text{number of } F, E \text{ pairs} \geq |D_\ell(B_i)| \left(\frac{\delta(\Lambda)}{2\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1} N_B}{|V(B)|} - 1 \right) \\ \therefore |D_\ell(B_i)| &\leq \frac{2^{k+2} \Delta(\Lambda)^2 |V(B)|}{\delta(\Lambda) K p_t^{\ell-1}}\end{aligned}$$

as $\frac{\delta(\Lambda)}{2\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1} N_B}{|V(B)|} - 1 \geq \frac{\delta(\Lambda)}{4\Delta(\Lambda)^2} \cdot \frac{Kp_t^{\ell-1} N_B}{|V(B)|}$ for sufficiently large K . Inserting this bound into (15), it is sufficient to show

$$\frac{Lq^2}{2n} |V(B)| - \sum_{\ell=1}^{k-1} \frac{2^{k+2} \Delta(\Lambda)^2 |V(B)|}{\delta(\Lambda) K p_t^{\ell-1}} \cdot \left(\frac{Lq^2}{n} \right)^\ell \geq \frac{2\Delta(\Lambda)}{\delta(\Lambda)} \cdot \frac{|V(B)|}{Kp_t^{k-1}} \left(\frac{Lq^2}{n} \right)^k$$

dividing through by $\frac{Lq^2|V(B)|}{n}$

$$\frac{1}{2} - \frac{2^{k+2}\Delta(\Lambda)^2}{\delta(\Lambda)K} \sum_{\ell=1}^{k-1} \left(\frac{Lq^2}{np_t}\right)^{\ell-1} \geq \frac{2\Delta(\Lambda)}{K\delta(\Lambda)} \left(\frac{Lq^2}{np_t}\right)^{k-1}$$

and so we want

$$K \geq \frac{2^{k+3}\Delta(\Lambda)^2}{\delta(\Lambda)} \sum_{\ell=0}^{k-1} \left(\frac{Lq^2}{np_t}\right)^{\ell} \tag{16}$$

Lastly, since we set $q = \frac{4c_P n}{|V(B)|}$, we have $\frac{Lq^2}{np_t} = \frac{16c_P^2 Ln}{p_t|V(B)|^2}$. And since B is given to be a rich block, we have

$$|V(B)|^2 \geq Lt^{2+\epsilon}|V(G)| = Ln / \left(\frac{n}{t^{2+\epsilon}|V(G)|}\right) = \frac{Ln}{p_t}$$

Thus, $\frac{Lq^2}{np_t} \leq 16c_P^2$ and we can set $K = \frac{2^{k+3}\Delta(\Lambda)^2}{\delta(\Lambda)} \cdot \frac{(16c_P^2)^{k-1}}{16c_P^2-1}$, which satisfies (16) and is not in terms of n , making it a valid definition for the constant K . \square

11 Conclusion

We have managed to show that the number of n -permutations that Λ -avoid π is bounded by $O\left(\frac{n \log^{2+\epsilon} n}{L}\right)^n$ only relying on the fact that Λ contains a certain collection of size- L cliques. This bound holds for positive ϵ arbitrarily close to 0. When L is polynomial in n , that is $L = \Theta(n^c)$ with $c \in (0, 1]$, this bound is a strict improvement on the $n! = O(n)^n$ total n -permutations. For $L = n$, we are a log-exponential factor off from the Stanley-Wilf conjecture.

Our matching bound for when Λ is a random hypergraph with edge probability α of $\exp(O(n))\alpha^{-\frac{n}{k-1}}$ is therefore more general in many ways, as there are no cliques of polynomial size in n w.h.p. in such a random graph. This is expected as the weakest part of our argument came from the deterministic nature of Λ . When we are bounding the sum of the vertices in the rich blocks,

$$\sum_{\text{rich blocks } B} |V(B)| = \sum_{\text{blocks } B} |V(B)| - \sum_{\text{unrich blocks } B} |V(B)|$$

the best bound for the unrich blocks

$$\sum_{\text{unrich blocks } B} |V(B)| \leq d(\text{number of unrich blocks}) \leq d(\text{number of blocks})$$

assumes that all the blocks are unrich, accounting for the worst deterministic case. When the locations of the blocks are randomized, we can make a stronger statement in expectation. However, such a reliance on large cliques in the fixed Λ case is necessary to achieve any meaningful bound, as we showed there are dense multipartite hypergraphs Λ^* which

are avoided by $O(n)^n$ permutations of length n , but which have constant maximal clique. This gives us hope that the conditions we place on the fixed Λ are relatively tight.

An open problem is to remove the $\log^{2+\epsilon} n$ term from the bound. The term comes from the use of hypergraph containers in a recursive branching fashion. Each container in the tree is broken down using the containers lemma as a black box, necessitating this term. It may be removable by reworking the arguments of the containers lemma to tailor to this recursive usage, which would improve our bound especially for $L = \Theta(n)$.

Acknowledgements

The authors would like to thank Asaf Ferber for his mentorship throughout this research. His teachings and advice were invaluable. They would also like to thank the reviewer for their helpful advice on improving the paper.

References

- [1] J. Balogh, R. Morris, and W. Samotij, “Independent sets in hypergraphs”, *J. Amer. Math. Soc.* 28, 669-709, 2015.
- [2] A. Ferber (personal communication, 2019).
- [3] A. Ferber, G. A. McKinley, and W. Samotij, “Supersaturated sparse graphs and hypergraphs”, *Int. Math. Research Notices*, 2020 (2), 378-402, 2020.
- [4] Z. Füredi and P. Hajnal, “Davenport-Schinzel theory of matrices”, *Discrete Mathematics*, 103 (3): 233-251, 1992.
- [5] M. Klazar, “The Füredi-Hajnal conjecture implies the Stanley-Wilf conjecture”, *Formal Power Series and Algebraic Combinatorics*, Springer, 250-255, Moscow, 2000.
- [6] A. Marcus and G. Tardos, “Excluded permutation matrices and the Stanley-Wilf conjecture”, *Journal of Combinatorial Theory, Series A*, 107 (1): 153-160, 2004.
- [7] R. Morris and D. Saxton, “The number of C_{2l} -free graphs”, *Adv. Math.* 298, 534-580, 2016.
- [8] D. Saxton and A. Thomason, “Hypergraph containers”, *Invent. Math.* 201, 925-992, 2015.