Extensions of the linear bound in the Füredi–Hajnal conjecture

Martin Klazar*   Adam Marcus†

May 17, 2006

Abstract

We present two extensions of the linear bound, due to Marcus and Tardos, on the number of 1-entries in an $n \times n$ (0,1)-matrix avoiding a fixed permutation matrix. We first extend the linear bound to hypergraphs with ordered vertex sets and, using previous results of Klazar, we prove an exponential bound on the number of hypergraphs on n vertices which avoid a fixed permutation. This, in turn, solves various conjectures of Klazar as well as a conjecture of Brändén and Mansour. We then extend the original Füredi–Hajnal problem from ordinary matrices to d-dimensional matrices and show that the number of 1-entries in a d-dimensional (0,1)-matrix with side length n which avoids a d-dimensional permutation matrix is $O(n^{d-1})$.

AMS subject classification: 05D05, 05C65.
keywords: Stanley-Wilf conjecture, extremal theory, (0,1)-matrix, ordered hypergraph.

1 Introduction

Füredi and Hajnal asked in [7] whether for every fixed permutation matrix P (that is, P has a single 1-entry in every row and column) the maximum

*Institute for Theoretical Computer Science (ITI) and Department of Applied Mathematics (KAM), Faculty of Mathematics and Physics of Charles University, Malostranské náměstí 25, 118 00 Praha, Czech Republic. Email: klazar@kam.mff.cuni.cz. ITI is supported by the project 1M0021620808 of the Czech Ministry of Education.

†Research performed as Visiting Researcher, Alfréd Rényi Institute, 1364 Budapest, Pf.127, Hungary, on leave from the Department of Mathematics (ACO), Georgia Institute of Technology, Atlanta, GA 30332-0160. Email: adam@math.gatech.edu. This research was made possible due to funding by the Fulbright Program in Hungary.

number of 1-entries in an $n \times n$ (0,1)-matrix M avoiding P is $O(n)$; the "avoidance" here means that P cannot be obtained from M by a series of row deletions, column deletions, and replacements of 1-entries with 0-entries (in particular, permuting rows or columns of M is not allowed). The Füredi–Hajnal conjecture was settled by Marcus and Tardos in [12] where they proved that if M avoids a $k \times k$ permutation matrix, then the number of 1-entries in M is at most $2k^4\binom{k^2}{k}n$. In this paper we present extensions of their linear bound to more general structures.

The Marcus–Tardos bound can be reformulated in the language of graph theory, since matrices with entries 0 and 1 can be viewed as the incidence matrices of bipartite graphs. In particular, if $P = ([2k], E(P))$ is a graph on the vertex set $[2k] = \{1, 2, \ldots, 2k\}$ with k mutually disjoint edges, each of which connects the sets $[k]$ and $[k+1, 2k] = \{k+1, k+2, \ldots, 2k\}$, and $M = ([2n], E(M))$ is a graph on $[2n]$ which only has edges connecting $[n]$ and $[n+1, 2n]$ and does not contain P as an ordered subgraph, then M has only linearly many edges, i.e. $|E(M)| = O(n)$. It is easy to modify the proof in [12] so that it gives a linear bound for all (not necessarily bipartite) P-avoiding graphs G on the vertex set $[2n]$ (and therefore $[n]$). In Section 2 we extend this bound further to hypergraphs with edges of arbitrary size. We also discuss exponential enumerative bounds which follow from the linear extremal bounds as corollaries.

In yet another light, (0,1)-matrices can be viewed as (the characteristic matrices of) binary relations. In Section 3 we generalize the original proof of Marcus and Tardos to d dimensions and show that every d-ary relation on $[n]$ which avoids a fixed d-dimensional permutation has at most $O(n^{d-1})$ elements (i.e., 1-entries in the characteristic matrix).

We conclude this section with a few remarks and references as to motivation for introducing ordered hypergraphs and their containment $\prec$. Besides the permutation avoidance touched upon in this article, the notion of "containment" that is studied here (defined below), first introduced by Klazar in [9], generalizes other well-studied situations in combinatorics: including the notion of noncrossing structures and Davenport–Schinzel sequences. Noncrossing structures S (i.e., graphs or set partitions) are those (ordered) structures having no four vertices $a < b < c < d$ such that a, c lie in one edge A and b, d lie in another edge B, where $A \neq B$; in the notation below, $S \not\succ (\{1, 3\}, \{2, 4\})$. Enumeration of noncrossing graphs goes back to the 19th century; see Flajolet and Noy [6] for more information and references. Noncrossing set partitions appear in many places in combinatorics and mathematics; see McCammond [13] and Simion [14] for surveys. The notion of noncrossing hypergraphs, see [9], is a common generalization of these two concepts, allowing edges of any size (as in partitions) as well as intersecting edges (as in graphs). The classical Davenport–Schinzel sequences are, in the framework of $\prec$, exactly those set partitions S satisfying $S \not\succ (\{1,3,5\},\{2,4\})$; see Klazar [10] and Valtr [16] for surveys. Thus $\prec$ provides a general framework to investigate extremal problems for ordered structures; see Tardos [15] for recent revival and progress in this area.

2 Extensions to hypergraphs

For a graph $G' = ([k], E')$, we define $\mathrm{gex}_<(n, G')$ to be the maximum number $|E|$ of edges in a graph $G = ([n], E)$ that does not contain G' as an ordered subgraph. We represent a permutation $\pi = a_1 a_2 \dots a_k$ of $[k]$ by the graph

$$P(\pi) = ([2k], \{\{i, k + a_i\} : \ i \in [k]\}).$$

As mentioned in Section 1, it is easy to modify the proof in [12] to obtain the bound

$$\mathrm{gex}_<(n, P(\pi)) = O(n) \tag{1}$$

where the hidden constant inside the O depends only on π.

For the hypergraph extension we need a few more definitions:

Definition. A *hypergraph* is a finite collection $\mathcal{H} = (E_i : \ i \in I)$ of finite nonempty edges E_i which are subsets of $\mathbf{N} = \{1, 2, \dots\}$. The *vertex set* is $V(\mathcal{H}) = \bigcup_{i \in I} E_i$. For simplicity we do not allow (unlike in the graph case) isolated vertices; for our extremal problems this restriction is immaterial (isolated vertices in graphs in this case can be represented by singleton edges). In general we *do* allow multiple edges, and will denote a hypergraph as *simple* if it has no multiple edges.

The *order* $v(\mathcal{H})$ of $\mathcal{H}$ is the number of vertices $v(\mathcal{H}) = |V(\mathcal{H})|$, the *size* $e(\mathcal{H})$ is the number of edges $e(\mathcal{H}) = |I|$, and the *weight* $i(\mathcal{H})$ is the number of incidences $i(\mathcal{H}) = \sum_{i \in I} |E_i|$.

We say that $\mathcal{H}' = (E_i' : \ i \in I')$ is *contained* in $\mathcal{H} = (E_i : \ i \in I)$, written $\mathcal{H}' \prec \mathcal{H}$, if there exists an increasing injection $f : V(\mathcal{H}') \to V(\mathcal{H})$ and an injection $g : I' \to I$ such that $f(E_i') \subset E_{g(i)}$ for every $i \in I'$; otherwise we say that $\mathcal{H}$ *avoids* $\mathcal{H}'$. That is, $\mathcal{H}' \prec \mathcal{H}$ means that $\mathcal{H}'$ can be obtained from $\mathcal{H}$ by deleting some edges, deleting vertices from the remaining edges, and relabeling the vertices so that their ordering is preserved.

This containment generalizes the ordered subgraph relation. Note that a simple hypergraph can contain a non-simple hypergraph. Throughout this article, we will use scripted letters to denote hypergraphs. Non-scripted letters will only be used when the structure must be a graph.

We define two hypergraph extremal functions.

Definition. Let $\mathcal{F}$ be any hypergraph. We associate with $\mathcal{F}$ the functions $\mathrm{ex}_e(\cdot, \mathcal{F}), \mathrm{ex}_i(\cdot, \mathcal{F}) : \mathbf{N} \to \mathbf{N}$,

$$\mathrm{ex}_e(n, \mathcal{F}) = \max\{e(\mathcal{H}) : \mathcal{H} \not\succ \mathcal{F} \ \& \ \mathcal{H} \text{ is simple} \ \& \ v(\mathcal{H}) \leq n\}$$
$$\mathrm{ex}_i(n, \mathcal{F}) = \max\{i(\mathcal{H}) : \mathcal{H} \not\succ \mathcal{F} \ \& \ \mathcal{H} \text{ is simple} \ \& \ v(\mathcal{H}) \leq n\}.$$

Obviously, $\mathrm{ex}_e(n, \mathcal{F}) \leq \mathrm{ex}_i(n, \mathcal{F})$ for every F and n.

If $A, B \subset \mathbf{N}$ and $a < b$ holds for every element $a \in A$ and $b \in B$, we write $A < B$. The sets A and B are said to be *separated*, if $A < B$ or $B < A$. If $\mathcal{F}$ has at least two edges and has no two separated edges, Theorem 2.3 in [11] gives an inequality in the opposite direction:

$$\mathrm{ex}_i(n, F) \leq (2v(F) - 1)(e(F) - 1)\mathrm{ex}_e(n, F)$$

So in particular, for every permutation π of $[k]$,

$$\mathrm{ex}_i(n, P(\pi)) \leq (4k - 1)(k - 1)\mathrm{ex}_e(n, P(\pi)). \tag{2}$$

Thus a linear bound on $\mathrm{ex}_i(n, P(\pi))$ follows directly from a linear bound on $\mathrm{ex}_e(n, P(\pi))$.

Such a bound on $\mathrm{ex}_e(n, P(\pi))$ can be derived using the techniques in [11] along with the graph bound in (1). To explain the reduction we need the notion of the *blow-up* of a graph.

Definition. Given a graph G, we say that the graph G' is an *m-blow-up of* G if for every edge coloring of G' by colors from $\mathbf{N}$ such that every color is used at most m times, there exists a subgraph of G' that is order-isomorphic to G and that has no two edges with the same color.

Let $G = (V, E)$ be a graph with k vertices and let B be a $\binom{k}{2}$-blow-up of G — for simplicity, we will use the graph B formed by replacing every vertex $v \in V$ with an independent set I_v of size $\binom{k}{2}$ and by placing all of the edges between I_v and I_u if and only if $\{u, v\} \in E$. Furthermore, let $f : \mathbf{N} \to \mathbf{N}$ be a function such that $\mathrm{gex}_<(n, B) < n \cdot f(n)$ for every n. Then by Theorem 3.1 in [11], we have that for all n,

$$\mathrm{ex}_e(n, G) < e(G) \cdot \mathrm{gex}_<(n, G) \cdot \mathrm{ex}_e(2f(n) + 1, G). \tag{3}$$

Combining the bounds in (1), (2), and (3) we obtain the following result:

Theorem 2.1. *For every permutation π,*

$$\mathrm{ex}_i(n, P(\pi)) = O(n).$$

Proof. For $m \in \mathbf{N}$ and a k-permutation π, we construct the permutation graph $P(\pi')$ from $P(\pi)$ by replacing every vertex of $P(\pi)$ with an interval of $k(m-1)+1$ vertices (so each $v \in [2k]$ becomes the interval $I_v = [(v-1)(km-k+1)+1, v(km-k+1)]$). Now for each edge $\{u, v\}$ in $P(\pi)$, we place a perfect matching between the intervals I_u and I_v. Thus if we take any selection of one edge from each of the k perfect matchings, the resulting graph is order-isomorphic to $P(\pi)$. It should be noted that there are many such $P(\pi')$'s (π' is always a $k^2(m-1)+k$-permutation) but each of them is, by the pigeonhole principle, an m-blow-up of $P(\pi)$.

We set $m = \binom{2k}{2}$. By the graph bound in (1), there are constants c_π and $c_{\pi'}$ such that

$$\mathrm{gex}_<(n, P(\pi)) < c_\pi n \quad \text{and } \mathrm{gex}_<(n, P(\pi')) < c_{\pi'} n$$

for every n. We set $B = P(\pi')$ and $f(n) = c_{\pi'}$ and apply the bound in (3) to get the linear bound

$$\mathrm{ex}_e(n, P(\pi)) < kc_\pi \cdot \mathrm{ex}_e(2c_{\pi'}+1, P(\pi)) \cdot n.$$

By the bound in (2),

$$\mathrm{ex}_i(n, P(\pi)) < k(k-1)(4k-1)c_\pi \cdot \mathrm{ex}_e(2c_{\pi'}+1, P(\pi)) \cdot n,$$

proving the claim. $\square$

Klazar posed the following six extremal and enumerative conjectures in [9]:

- C1: The number of simple $\mathcal{H}$ such that $v(\mathcal{H}) = n$ and $\mathcal{H} \not\succ P(\pi)$ is at most c_1^n.
- C2: The number of maximal simple $\mathcal{H}$ with $v(H) = n$ and $\mathcal{H} \not\succ P(\pi)$ is at most c_2^n.
- C3: For every simple $\mathcal{H}$ with $v(\mathcal{H}) = n$ and $\mathcal{H} \not\succ P(\pi)$, we have $e(\mathcal{H}) < c_3 n$.
- C4: For every simple $\mathcal{H}$ with $v(\mathcal{H}) = n$ and $\mathcal{H} \not\succ P(\pi)$, we have $i(\mathcal{H}) < c_4 n$.
- C5: The number of simple $\mathcal{H}$ with $i(\mathcal{H}) = n$ and $\mathcal{H} \not\succ P(\pi)$ is at most c_5^n.

C6: The number of $\mathcal{H}$ with $i(\mathcal{H}) = n$ and $\mathcal{H} \not\succ P(\pi)$ is at most c_6^n.

He showed that all six conjectures hold for a large class of permutations π and that they hold for every π in weaker forms: with almost linear and almost exponential bounds (respectively). Conjecture C4, however, is precisely the statement of Theorem 2.1, and it is easy to extend the proof given in this paper to affirm that all six conjectures hold for every permutation π.

We shall show how to amend the proofs in [9] to prove C1, and then note that C1 implies C2, C3, C5 and C6 via Lemma 2.1 of [9].

Corollary 2.2. *For every permutation π there exists a constant $c_1 > 1$ (depending on π) so that the number of simple hypergraphs on the vertex set $[n]$ avoiding $P(\pi)$ is less than c_1^n.*

Proof. Theorems 2.4 and 2.5 in [9] show that the number of hypergraphs with a given weight $i(H)$ that avoid $P(\pi)$ is at most $9^{(3^{2k}+2k)i(H)}$. Thus by Theorem 2.1, we are done. □

The Stanley–Wilf conjecture (see Bóna [2]), proved by Marcus and Tardos in [12] as a corollary of their linear extremal bound, claimed that for every permutation π there is a constant $c = c(\pi)$ such that the number of permutations σ of $[n]$ avoiding π is less than c^n; the avoidance of permutations here means that $P(\pi)$ is not an ordered subgraph of $P(\sigma)$. In view of the reformulation from permutations to bipartite graphs mentioned in Section 1, Corollary 2.2 is an extension of the Stanley–Wilf conjecture. A related extension was proposed by Brändén and Mansour in Section 5 of [4]: they conjectured that the number of *words* over the ordered alphabet $[n]$ which have length n and avoid π is at most exponential in n. These words can be represented by simple graphs G on $[2n]$ in which every edge connects $[n]$ and $[n+1, 2n]$ and every $x \in [n]$ has degree exactly 1; the containment of ordered words is then just the ordered subgraph relation. Hence this extension is subsumed in Corollary 2.2.

Corollary 2.2 subsumes yet another extension of the Stanley–Wilf conjecture to set partitions proposed by Klazar [8]. This extension is related to *k-noncrossing* and *k-nonnesting* set partitions whose exact enumeration was recently investigated by Chen et al. [5] and Bousquet-Mélou and Xin [3]. Consider, for a set partition S of $[n]$, the graph $G(S) = ([n], E)$ in which an edge connects two neighboring elements of a block (not separated by another element of the same block). Then S is represented by increasing paths which are spanned by the blocks. S is a k-noncrossing (resp. k-nonnesting) partition if and only if $P(12\ldots k)$ (resp. $P(k(k-1)\ldots 1)$) is not an ordered subgraph of $G(H)$. Thus Corollary 2.2 provides an exponential bound: for fixed k, the numbers of k-noncrossing and k-nonnesting partitions of $[n]$ grow at most exponentially.

3 An extension to d-dimensional matrices

We now generalize the original Füredi–Hajnal conjecture from ordinary (0,1)-matrices to d-dimensional (0,1)-matrices. As was mentioned in Section 1, these are just d-ary relations (or, in hypergraph terminology, d-uniform, d-partite hypergraphs). We keep the matrix terminology, however, both for the sake of consistency and to highlight the similarities with the original Marcus–Tardos proof in [12].

Definition. We will call a $(d+1)$-tuple $\mathcal{M} = (M; n_1, \ldots, n_d)$ where $M \subset [n_1] \times \cdots \times [n_d]$ a *d-dimensional (0,1)-matrix*, and will refer to the elements of M as *1-entries* of M. We define the *size* of $\mathcal{M}$ (written $|\mathcal{M}|$) to be the cardinality of the set M (the number of 1-entries).

If $\mathcal{F} = (F; k_1, \ldots, k_d)$ and $\mathcal{M} = (M; n_1, \ldots, n_d)$ are two d-dimensional matrices, we say that F is *contained* in $\mathcal{M}$, written $\mathcal{F} \prec \mathcal{M}$, if there exist d increasing injections $f_i : [k_i] \to [n_i]$, $i = 1, 2, \ldots, d$, such that for every $(x_1, \ldots, x_d) \in F$ we have $(f_1(x_1), \ldots, f_d(x_d)) \in M$; otherwise we say that $\mathcal{M}$ *avoids* $\mathcal{F}$.

Note that in hypergraph terminology, $\mathcal{M} = (M; n_1, \ldots, n_d)$ is nothing more than a d-partite, d-uniform (ordered) hypergraph with the i^{th} partition having n_i vertices. Then M would be the collection of edges (where $(x_1, x_2, \ldots, x_d)$ would be a 1-entry in $\mathcal{M}$ if and only if $\{x_1, x_2, \ldots, x_d\}$ was an edge in the hypergraph).

Definition. We set $f(n, \mathcal{F}, d)$ to be the largest size $|\mathcal{M}|$ of a d-dimensional matrix $\mathcal{M} = (M; n, \ldots, n)$ that avoids a fixed d-dimensional matrix F.

For $i \in [d]$, we denote the *projection map* from $[n_1] \times \cdots \times [n_d]$ to $[n_i]$ as ρ_i.

For $t \in [d]$, we define the *t-remainder* of $\mathcal{M} = (M; n_1, \ldots, n_d)$ to be the $(d-1)$-dimensional matrix $\mathcal{N} = (N; n'_1, \ldots, n'_{d-1})$ where $n'_1 = n_1, \ldots, n'_{t-1} = n_{t-1}$, $n'_t = n_{t+1}, \ldots, n'_{d-1} = n_d$ and the entry $(e_1, \ldots, e_{d-1}) \in N$ if and only if $(e_1, \ldots, e_{t-1}, x, e_t, e_{t+1}, \ldots, e_{d-1}) \in M$ for some $x \in [n_t]$.

Let $I_1 < I_2 < \cdots < I_r$ be a partition of $[n]$ into r intervals and $\mathcal{M} = (M; n, \ldots, n)$ a d-dimensional matrix. We define the *contraction* of $\mathcal{M}$ (with respect to the intervals) to be the d-dimensional matrix $\mathcal{N} = (N; r, \ldots, r)$ given by $(e_1, \ldots, e_d) \in N$ if and only if $M \cap (I_{e_1} \times \cdots \times I_{e_d}) \neq \emptyset$ (we could define the contraction operation for a general d-dimensional matrix and with distinct and general partitions in each coordinate but we will not need such generality).

Again, translating into the hypergraph terminology, the image of M by the projection ρ_i is obtained by intersecting the edges (i.e. 1-entries) with the i^{th} part in the partition, while the intersections with the union of all parts except the t^{th} one gives the t-remainder of M (in both cases we disregard multiplicity of edges). Also, the contraction of $\mathcal{M}$ with respect to a partition I can be viewed as a block hypergraph of $\mathcal{M}$. That is, the edge $\{x_1, x_2, \ldots, x_d\}$ exists in the contraction if and only if there is at least one edge in $\mathcal{M}$ which contains a vertex from *each* of the blocks $B_{x_1}, B_{x_2}, \ldots, B_{x_n}$. Here, the blocks are the parts of the underlying vertex partition each cut into pieces by the partition I.

Definition. For $d \geq 2$, we say that $\mathcal{P} = (P; k, \ldots, k)$ is a *d-dimensional permutation of* $[k]$ if for every $i \in [d]$ and $x \in [k]$ there is a single 1-entry $e \in P$ with $\rho_i(e) = x$. In the degenerate case $d = 1$, we define the only 1-dimensional permutation to be $\mathcal{P} = (P; k)$ with $P = [k]$. Note that there are exactly $(k!)^{d-1}$ different d-dimensional permutations of $[k]$, and that each has exactly k 1-entries. For example, the 2-dimensional permutations $\mathcal{P} = (P; k, k)$ are precisely the $k \times k$ permutation matrices.

In hypergraph terminology, the d-dimensional permutations of $[k]$ would be the set of perfect matchings of the complete d-uniform, d-partite hypergraph on kd vertices.

We will make use of two observations, analogous to those made in [12]:

1. For any $t \in [d]$, the t-remainder of a d-dimensional permutation of $[k]$ is a $(d-1)$-dimensional permutation of $[k]$. Furthermore, each 1-entry of the resulting t-remainder can be completed (by adding the t-th coordinate) in a unique way to an edge of the original permutation.

2. If $\mathcal{M} = (M; n, \ldots, n)$ avoids a d-dimensional permutation, then so does any contraction of $\mathcal{M}$.

Our goal is to prove the following theorem, which is a generalization of the result due to Marcus and Tardos:

Theorem 3.1. *For every fixed d-dimensional permutation $\mathcal{P}$,*

$$f(n, \mathcal{P}, d) = O(n^{d-1}).$$

On the other hand it is clear that for a d-dimensional permutation $\mathcal{P}$ with $|\mathcal{P}| > 1$ we have $f(n, \mathcal{P}, d) \geq n^{d-1}$ ($f(n, \mathcal{P}, d) = 0$ if $|\mathcal{P}| = 1$; $|\mathcal{P}| = 1$ can only occur when $k = 1$). Thus, for a d-dimensional permutation $\mathcal{P}$ with $|\mathcal{P}| > 1$,

$$f(n, \mathcal{P}, d) = \Theta(n^{d-1}).$$

To prove Theorem 3.1, we will show that a d-dimensional matrix of big enough size must contain every d-dimensional permutation of k. We set

$$f(n, k, d) = \max_{\mathcal{P}} f(n, \mathcal{P}, d)$$

where $\mathcal{P}$ runs through all d-dimensional permutations of $[k]$.

Lemma 3.2. *Let $d \geq 2$, $m, n_0 \in \mathbf{N}$. Then*

$$f(mn_0, k, d) \leq (k-1)^d \cdot f(n_0, k, d) + dn_0 m^d \binom{m}{k} \cdot f(n_0, k, d-1).$$

Proof. Let $\mathcal{M} = (M, mn_0, \ldots, mn_0)$ be a d-dimensional matrix that avoids $\mathcal{P}$, a d-dimensional permutation of $[k]$. We aim to bound the size of $\mathcal{M}$.

We split $[mn_0]$ into n_0 intervals $I_1 < I_2 < \cdots < I_{n_0}$, each of length m, and define, for $i_1, \ldots, i_d \in [n_0]$,

$$S(i_1, \ldots, i_d) = \{e \in M : \rho_j(e) \in I_{i_j} \text{ for } j = 1, \ldots, d\}.$$

Note that this partitions the set of 1-entries of $\mathcal{M}$ into n_0^d (possibly empty) collections. We will call these collections *blocks* and we define a cover of the blocks by a total of $dn_0 + 1$ sets $\{U_0\} \cup \{U(t, j) : t \in [d], j \in [n_0]\}$ as follows:

- $U(t, j) = \{S(i_1, \ldots, i_d) : \ i_t = j \text{ and } |\rho_t(S(i_1, \ldots, i_d))| \geq k\}$;
- U_0 consists of the blocks which are not in any $U(t, j)$.

Note that the total number of non-empty blocks is exactly the number of 1-entries in the contraction of $\mathcal{M}$ with respect to the partition $\{I_i\}$. Since $\mathcal{M}$ does not contain $\mathcal{P}$, the contraction of $\mathcal{M}$ cannot contain $\mathcal{P}$, so the number of non-empty blocks is at most $f(n_0, k, d)$. Also note that any block B in U_0 has at most $(k-1)^d$ non-zero entries in it (because $B \subset X_1 \times \cdots \times X_d$ for some $X_i \subset [mn_0]$ with $|X_i| < k$). Hence

$$|\bigcup U_0| \leq (k-1)^d \cdot f(n_0, k, d).$$

Now we fix $t \in [d]$ and $j \in [n_0]$. Clearly,

$$|\bigcup U(t, j)| \leq m^d |U(t, j)|.$$

We assume, for a contradiction, that $|U(t,j)| > \binom{m}{k} \cdot f(n_0, k, d-1)$. By the definition of $U(t,j)$ and the pigeonhole principle, there are k numbers $c_1 < c_2 < \cdots < c_k$ in I_j and r blocks $S_1, S_2, \ldots, S_r$ in $U(t,j)$ where $r > f(n_0, k, d-1)$ such that for every S_a and every c_b there is an $e \in S_a$ with $\rho_t(e) = c_b$. Let $\mathcal{P}'$ be the t-remainder of $\mathcal{P}$ and $\mathcal{M}' = (M'; n_0, \ldots, n_0)$ be the $(d-1)$-dimensional matrix arising from contracting $(\bigcup_{i=1}^r S_i, n, \ldots, n)$ with respect to the intervals I_i and then taking the t-remainder. Since $|\mathcal{M}'| = r > f(n_0, k, d-1)$, $\mathcal{M}'$ contains $\mathcal{P}'$. Thus among the blocks $S_1, S_2, \ldots, S_r$ there exist k of them — call them $S_1, S_2, \ldots, S_k$ — so that for any selection of k edges $e_1 \in S_1, \ldots, e_k \in S_k$ their t-remainders form a copy of $\mathcal{P}'$. Furthermore, due to the property of the blocks S_i, it is possible to select $e_1, \ldots, e_k$ so that their t-th coordinates agree with $\mathcal{P}$. Then $e_1, \ldots, e_k$ form a copy of $\mathcal{P}$, a contradiction. Therefore

$$|\bigcup U(t,j)| \leq m^d |U(t,j)| \leq m^d \tbinom{m}{k} \cdot f(n_0, k, d-1)$$

and

$$|\bigcup_{t,j} \bigcup U(t,j)| \leq d n_0 m^d \tbinom{m}{k} \cdot f(n_0, k, d-1).$$

Combining this with the bound for U_0 gives the stated bound. □

Theorem 3.1 will be a direct consequence of the following lemma:

Lemma 3.3. *If $m = \lceil k^{d/(d-1)} \rceil$, then $f(n,k,d) \leq k^d \big(dm\binom{m+1}{k}\big)^{d-1} n^{d-1}$.*

Proof. We will proceed by induction on $d + n$. For $d = 1$ this holds since $f(n,k,1) = k-1$ and, for $n < k^2$, this holds trivially. Now given n and $d \geq 2$, assume that the hypothesis is true for all d', n' such that $d' + n' < d + n$.

Let $n_0 = \lfloor n/m \rfloor$ and

$$c_d = k^d \left(dm \binom{m+1}{k} \right)^{d-1}.$$

Using the inequality $f(n,k,d) < f(mn_0, k, d) + dmn^{d-1}$, Lemma 3.2, the inductive hypotheses, and $n_0 \leq n/m$, we get

$$f(n,k,d) < \left(\frac{(k-1)^d}{m^{d-1}} c_d + dm \left(\binom{m}{k} c_{d-1} + 1 \right) \right) n^{d-1}.$$

Since $\frac{(k-1)^d}{m^{d-1}} \leq (1 - \frac{1}{k})^d \leq 1 - \frac{1}{k}$ and $\binom{m}{k} c_{d-1} + 1 \leq \binom{m+1}{k} c_{d-1}$, it follows that $f(n,k,d) < c_d n^{d-1}$ with the above defined c_d. □

4 Concluding remarks

We were informed recently that Balogh, Bollobás and Morris [1] derived Theorem 2.1 (their Theorem 2) and Corollary 2.2 (their Theorem 1) independently. The proofs in [1] are self-contained (not appealing to the results in [11]) and their approach is different from ours. In fact, they are able to prove stronger statements, which in turn imply Theorem 2.1 and Corollary 2.2 from this paper.

It should be noted that we make no effort to optimize any of the constants in Section 3, however it would be interesting to see if any of the constants could be drastically reduced. The constant achieved in this paper is double exponential in k, whereas we suspect that the true constant is in fact much smaller.

5 Acknowledgments

The authors would like to thank Gábor Tardos for, among other things, his enlightening conversations and helpful remarks, and the referee for careful reading of the manuscript.

References

[1] J. Balogh, B. Bollobás and R. Morris, Hereditary properties of partitions, ordered graphs and ordered hypergraphs, European J. Combin., to appear.

[2] M. Bóna, Combinatorics of Permutations, Chapman & Hall/CRC, Boca Raton, FL, USA 2004.

[3] M. Bousquet-Mélou and G. Xin, On partitions avoiding 3-crossings, Sém. Lothar. 54 (2005) B54e, 21 pages.

[4] P. Brändén and T. Mansour, Finite automata and pattern avoidance in words, J. Combin. Theory Ser. A 110 (2005) 127–145.

[5] W. Y. C. Chen, E. Y. P. Deng, R. R. X. Du, R. P. Stanley and C. H. Yan, Crossings and nestings of matchings and partitions, Trans. Amer. Math. Soc., to appear.

[6] P. Flajolet and M. Noy, Analytic combinatorics of non-crossing configurations, Discrete Math. 204 (1999) 203–229.

[7] Z. Füredi and P. Hajnal, Davenport–Schinzel theory of matrices, Discrete Math. 103 (1992) 233–251.

[8] M. Klazar, Counting pattern-free set partitions. I. A generalization of Stirling numbers of the second kind, European J. Combin. 21 (2000) 367–378.

[9] M. Klazar, Counting pattern-free set partitions. II. Noncrossing and other hypergraphs, Electron. J. Combin. 7 (2000), research paper 34, 25 pp.

[10] M. Klazar, Generalized Davenport-Schinzel sequences: results, problems, and applications, Integers 2 (A11) 2002, 39 p.

[11] M. Klazar, Extremal problems for ordered (hyper) graphs: applications of Davenport–Schinzel sequences, European J. Combin. 25 (2004) 125–140.

[12] A. Marcus and G. Tardos, Excluded permutation matrices and the Stanley–Wilf conjecture, J. Combin. Theory Ser. A 107 (2004) 153–160.

[13] J. McCammond, Noncrossing partitions in surprising locations, Amer. Math. Monthly, to appear

[14] R. Simion, Noncrossing partitions, Discrete Math. 217 (2000) 367–409.

[15] G. Tardos, On 0-1 matrices and small excluded submatrices, J. Combin. Theory Ser. A 111 (2005) 266–288.

[16] P. Valtr, Generalizations of Davenport-Schinzel sequences, in: R. L. Graham, et al., (Eds.), Contemporary Trends in Discrete Mathematics, American Mathematical Society, Providence, RI, 1999, pp. 349–389.