

Forbidden Submatrices: Some new bounds and constructions

R. P. Anstee* and Ruiyuan Chen†

Mathematics Department
The University of British Columbia
Vancouver, B.C. Canada V6T 1Z2

anstee@math.ubc.ca ruiyuan@interchange.ubc.ca

Submitted: Mar 24, 2012; Accepted: Jan 2, 2013; Published: Jan 14, 2013

Abstract

We explore an extremal hypergraph problem for which both the vertices and edges are ordered. Given a hypergraph F (not necessarily simple), we consider how many edges a simple hypergraph (no repeated edges) on m vertices can have while forbidding F as a subhypergraph where both hypergraphs have fixed vertex and edge orderings. A hypergraph of n edges on m vertices can be encoded as an $m \times n$ (0,1)-matrix. We say a matrix is *simple* if it is a (0,1)-matrix with no repeated columns. Given a (0,1)-matrix F , we define $\text{fs}(m, F)$ as the maximum, over all simple matrices A which do not have F as a submatrix, of the number of columns in A . The row and column order matter. It is known that if F is $k \times \ell$ then $\text{fs}(m, F)$ is $O(m^{2k-1-\epsilon})$ where $\epsilon = (k-1)/(13 \log_2 \ell)$. Anstee, Frankl, Füredi and Pach have conjectured that if F is k -rowed, then $\text{fs}(m, F)$ is $O(m^k)$. We show $\text{fs}(m, F)$ is $O(m^2)$ for $F = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & \cdots \\ 0 & 1 & 0 & 1 & 0 & \cdots \end{bmatrix}$ and for $F = \begin{bmatrix} 1 & 0 & 1 & 0 & 1 & \cdots \\ 1 & 0 & 1 & 0 & 1 & \cdots \end{bmatrix}$. The proofs use a type of amortized analysis. We also give some constructions.

Keywords: extremal set theory, forbidden submatrix, ordered sets, trace, amortized analysis

1 Introduction

We are considering a problem in extremal set theory. We find it convenient to use the language of matrix theory and of sets. Let $[m] = \{1, 2, \dots, m\}$. We define a *simple* matrix as a (0,1)-matrix with no repeated columns. An $m \times n$ simple matrix A can be thought

*Research supported in part by NSERC

†Research supported by NSERC USRA and NSERC grant of first author

of as a simple hypergraph $\mathcal{A}$ of n edges $S_1, S_2, \dots, S_n$, each a subset of $[m]$, where $i \in S_j$ if and only if the i, j entry of A is 1.

We are interested in the property of forbidding a *submatrix* F . Here we are concerned about both row and column order. This would be a subhypergraph with vertex and edges ordered. The 1's are as important as the 0's unlike in the study of *patterns* [7] where 1's are the dominant focus. We could forbid the *pattern* associated with the 1's of F by forbidding as submatrices, all matrices G such that $F \leq G$ i.e. all matrices with 1's in the same position as F . Typically the goal of forbidden patterns is an upper bound on the number of 1's.

Note that forbidding a submatrix F is a weaker condition than the condition of forbidding all row and column permutations of F as a submatrix. This latter notion is that of a *forbidden configuration* studied in various papers including [2].

Sometimes we are able to obtain structural results about a simple matrix A with no submatrix F . A case of some interest for optimization problems, considered in [6] and [3], were *totally balanced* matrices (defined by a family of forbidden configurations) which were precisely those matrices for which we could order the rows and columns so that the resulting matrix had no submatrix

$$\Gamma = \begin{bmatrix} 1 & 1 \\ 1 & 0 \end{bmatrix}.$$

Γ -free matrices are explored in a number of papers.

Our goal in this paper is to consider the extremal problem of how many columns an m -rowed simple matrix A can have under a forbidden submatrix property. Let $\|A\|$ denote the number of columns of a matrix A . Let F be given. We define

$$\text{fs}(m, F) = \max_A \{\|A\| : A \text{ is } m\text{-rowed simple with no submatrix } F\}.$$

Alternatively $\text{fs}(m, F)$ is the smallest value so that if A is an $m \times (\text{fs}(m, F) + 1)$ simple matrix then A must contain the submatrix F . Note that we do not require F to be simple in the definition; most forbidden submatrices in this paper are non-simple. The best general result is the following.

Theorem 1.1 [1] *Let F be a $k \times \ell$ $(0,1)$ -matrix. Then $\text{fs}(m, F)$ is $O(m^{2k-1-\epsilon})$ where $\epsilon = (k-1)/(13 \log_2 \ell)$.*

There is a conjecture on the asymptotics of $\text{fs}(m, F)$.

Conjecture 1.2 [4],[5] *Anstee, Frankl, Füredi, Pach. Let F be a given $k \times \ell$ $(0,1)$ -matrix. Then $\text{fs}(m, F)$ is $O(m^k)$.*

The evidence for the conjecture is not extensive. We do know the conjecture is true for $1 \times \ell$ F [4] and also for the $k \times 2$ matrix of a column of 1's followed by a column of 0's [5] as well as a few other cases. The conjecture has not been established for $2 \times \ell$ F . Theorems 1.4 and 1.5 are modest progress and further evidence for the conjecture while showing the utility of an amortized approach. We have been unable to extend the amortized approach to general $2 \times \ell$ F .

Lemma 1.3 [4] *Let α be a $k \times 1$ $(0,1)$ -column. Then*

$$fs(m, [\alpha]) = \binom{m}{k-1} + \binom{m}{k-2} + \cdots + \binom{m}{0}$$

i.e. $fs(m, \alpha)$ is $\Theta(m^{k-1})$.

Proof: Let $\alpha = (a_1, a_2, \dots, a_k)^T$. Consider any $m \times 1$ $(0,1)$ -column β not containing α . Some notation can help in describing the columns which we will treat as strings of 0's and 1's. For $a \in \{0, 1\}$, let $\bar{a}$ denote the $(0,1)$ -complement of a , and let a^* denote a string of zero or more copies of a . Then for each β there will be a j with $0 \leq j < k$ so that column β will contain $(a_1, a_2, \dots, a_j)^T$ as a submatrix and not contain $(a_1, a_2, \dots, a_{j+1})^T$. Then β is some instance of

$$((\bar{a}_1)^* a_1 (\bar{a}_2)^* a_2 (\bar{a}_3)^* \cdots a_j (\bar{a}_{j+1})^*)^T.$$

The number of columns β which contain $(a_1, a_2, \dots, a_j)^T$ and not $(a_1, a_2, \dots, a_{j+1})^T$ is seen to be $\binom{m}{j}$ considering the choices for the positions of $a_1, a_2, \dots, a_j$ in β . Summing over all choices for j we obtain $fs(m, \alpha) = \binom{m}{k-1} + \binom{m}{k-2} + \cdots + \binom{m}{0}$. ■

It is interesting that Conjecture 1.2 proposes such a narrow range on the asymptotics of $fs(m, F)$, since for k -rowed F by Lemma 1.3, $fs(m, F)$ is $\Omega(m^{k-1})$. Moreover for F containing two non-identical columns, then by Proposition 5.4, there is a $\Omega(m^k)$ construction avoiding F . Several results which confirm Conjecture 1.2 are in [4] and [5]. The evidence is limited. We are able to offer some additional evidence and some constructions avoiding certain F that may be useful in considering this problem.

Theorem 1.4 *Let ℓ be given. Let F be the $2 \times \ell$ matrix*

$$F = \begin{bmatrix} 1 & 0 & 1 & 0 & \cdots \\ 0 & 1 & 0 & 1 & \cdots \end{bmatrix}. \quad (1)$$

Then

$$(\ell - 1) \binom{m}{2} - \ell^2 m - 3\ell^2 \leq fs(m, F) \leq (\ell - 1) \binom{m}{2} + m + 1. \quad (2)$$

Theorem 1.5 *Let ℓ be given. Let F be the $2 \times \ell$ matrix*

$$F = \begin{bmatrix} 1 & 0 & 1 & 0 & \cdots \\ 1 & 0 & 1 & 0 & \cdots \end{bmatrix}. \quad (3)$$

Then

$$(\ell - 1) \binom{m}{2} - \ell^2 m - 3\ell^2 \leq fs(m, F) \leq 6(\ell - 1) \binom{m}{2} + m + 2. \quad (4)$$

This does not answer Conjecture 1.2 but may stimulate further work. The bound in Theorem 1.4 for $\ell = 2$ can be found in [4] while the exact bound for Theorem 1.5 for $\ell = 2$ of $\binom{m}{2} + 2m - 1$ is established in [5]. Also note that a row interchange of (1) contains a $2 \times (\ell - 1)$ submatrix (1) and the matrix (3) is not affected by row permutations. Thus row order is not very important in these cases. We introduce some constructions in Proposition 5.1, Proposition 5.4 in Section 5. In Corollary 5.2, we use the following easy monotonicity result but point out that the corresponding result for a single forbidden configuration is not known.

Lemma 1.6 *For a given F and $m \geq 2$, we have $fs(m, F) \geq fs(m - 1, F)$.*

Proof: Let A be an $(m - 1) \times fs(m - 1, F)$ simple matrix with no submatrix F . Then either $\begin{bmatrix} A \\ 00\dots 0 \end{bmatrix}$ or $\begin{bmatrix} A \\ 11\dots 1 \end{bmatrix}$ is an $m \times fs(m - 1, F)$ simple matrix with no submatrix F . ■

It is known that there are families of forbidden configurations not having monotonicity as above. For forbidden families of submatrices there will also be problems. A simple example is to forbid all 2^k $k \times 1$ (0,1)-submatrices for which the bound would be 2^{k-1} for $m = k - 1$ and the bound would be 0 for $m \geq k$.

For a subset $S \subseteq [m]$ of rows, let $A|_S$ denote the submatrix of A given by rows S . Let $[A|B]$ denote the concatenation of two matrices A, B on the same number of rows. We define the *lower triangular* matrix T_m as the $m \times m$ (0,1)-matrix with 1's in positions i, j if and only if $i \geq j$. Let $\mathbf{0}_m$ denote the $m \times 1$ column of all 0's.

2 Amortized analysis

Let F be a k -rowed matrix. In order to search for a submatrix F in a matrix A , the natural idea is to process the columns of A from left to right keeping track of how much each k -set of rows already contains of the initial columns of F . For each k -set, we have a *bucket*. Each bucket will contain a k -rowed submatrix of A , possibly of no columns. Let an *initial submatrix* of F be the submatrix formed by taking the first i columns for some i . After processing say the first t columns of A , corresponding to the $m \times t$ submatrix $A(t)$ of A of its first t columns, the bucket associated with a k -set S will hold the largest initial submatrix of F contained in $A(t)|_S$.

A *contribution* is an addition of a column to a bucket. We refer to a column of A as a *contributing* column if it adds a column to at least one bucket. We may think of our buckets as a pushdown stack, the top of the stack recording the last contribution. The number of contributions of a contributing column is the number of buckets that are augmented. A *filler* column is one that makes no contribution. It can be shown that we can process at most $O(m^{k-1})$ filler columns before encountering a contributing column. After $(\ell - 1)\binom{m}{k} + 1$ contributions at least one bucket will have a submatrix of ℓ columns yielding the forbidden submatrix F . This yields that after $O(m^{2k-1})$ columns A has a submatrix F from the pigeonhole argument as reported in [5]. We need a more detailed analysis. An amortized approach was used to handle the 1-rowed $F = [1010\dots]$ in [1].

In our proofs of Theorem 1.4 and Theorem 1.5, we partition A into blocks $A = [A_0|A_1|A_2|\cdots]$, where each block A_i (except A_0) consists of a single contributing column α_i followed by any number of filler columns. Block A_0 has no contributing columns and hence may have no columns. Consider a contributing column with t contributions. Then we imagine that we have credits (t credits in our proof of Theorem 1.4 and $6t$ credits in our proof of Theorem 1.5) sufficient to pay for the filler columns with a limited number of exceptions ($m+1$ in Theorem 1.4 and $m+2$ in Theorem 1.5). This will yield the results.

The following two lemmas give us some useful structural information for a block A_i . The first considers matrix (1).

Lemma 2.1 *Assume that an m -rowed simple matrix A has neither submatrix $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ nor $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Then A is a submatrix of some row and column permutation of $[T_m|\mathbf{0}_m]$.*

Note that any pair of columns of A of the same column sum will contain either $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ or $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$. Moreover a column α of sum r and a column β of sum s with $r < s$ cannot have row where α is 1 and β is 0 since then there will also be a row where α is 0 and β is 1 yielding either $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ or $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$.

The second lemma considers matrix (3). Define a m -rowed *sunflower* A as an m -rowed matrix which has a *centre* $X \subseteq [m]$ and columns which are all 1's on rows X and have at most one 1 in a row of $[m] \setminus X$. The sunflower idea is well used in [5]. We define the *centre column* $\psi(A)$ of the sunflower as the column which is all 1's on rows X and 0's on $[m] \setminus X$ but note that this column need not be a column of A . We also consider the $(0,1)$ -complement of A , which we call a *inverse sunflower* whose centre is X (all columns in the inverse sunflower will be all 0's on the rows X).

Lemma 2.2 *Assume that an m -rowed simple matrix A has neither submatrix $\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}$ nor $\begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$. Then the columns of A form an m -rowed sunflower or inverse sunflower.*

The proof is easy for A having one or two columns. Once A has 3 columns, we may uniquely determine the centre X and whether A is a sunflower or inverse sunflower. Any additional columns will simply extend the sunflower or inverse sunflower [5].

3 Checkerboard

Proof of Theorem 1.4: Let F be the matrix of (1). We may use Corollary 5.3 to establish the lower bound for $\text{fs}(m, F)$. For the upper bound, partition A into blocks $A = [A_0|A_1|A_2|\cdots]$, where each block A_i (except A_0) consists of a single contributing column α_i followed by any number of filler columns. A_0 has no contributing columns and hence may have no columns. We have that A_i has neither submatrix $\begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix}$ nor $\begin{bmatrix} 0 & 1 \\ 1 & 0 \end{bmatrix}$ and so by Lemma 2.1, the columns of A_i are contained in some row permutation of $[T_m|\mathbf{0}_m]$. We will show a little more below.

We assign one credit to each contribution of a contributing column and then show that all columns (except possibly up to $m+1$ columns from $[T_m|\mathbf{0}_m]$ which are all possible filler columns initially) in A will be paid for by credits. This yields the bound (2).

We can record the state of our $\binom{m}{2}$ buckets as follows. Form a digraph $D = ([m], E)$ where we have a directed edge $p \rightarrow q \in E$ if we are looking for $\begin{smallmatrix} p \\ q \end{smallmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ to add to our bucket $\{p, q\}$. What we are looking for on a pair of rows is the opposite of the top of the stack in the bucket of that pair of rows. If the bucket is empty (all buckets will be empty initially) then for $p < q$ we are looking for $\begin{smallmatrix} p \\ q \end{smallmatrix} \begin{bmatrix} 1 \\ 0 \end{bmatrix}$ and so have the directed edge $p \rightarrow q$. We necessarily have a *tournament* (for each bucket $\{p, q\}$, we have $p \rightarrow q$ or $q \rightarrow p$) and initially we have a *transitive tournament* ($p \rightarrow q$ and $q \rightarrow r$ implies $p \rightarrow r$) or what can be called a *total order* of the rows $1, 2, \dots, m$ so that for all pairs p, q with $p < q$ in that order, we have $p \rightarrow q$. After we process a contributing column the directed graph changes, for each contribution to a bucket $\{u, v\}$ we reverse the arc in D joining u and v . We will show by induction that when we process a contributing column that the resulting digraph is still a total order. Assume that after processing α_{i-1} , we have a total order that is in fact $1, 2, \dots, m$. We would typically have to permute the rows of A for this to be true, but note that this is the correct order in the case $i = 1$ before processing any contributing column. Now consider processing the next contributing column $\alpha_i = (a_1, a_2, \dots, a_m)^T$. Let $b_1 < b_2 < \dots < b_s$ be the rows containing the 0's of α_i and let $c_1 < c_2 < \dots < c_t$ be the rows containing the 1's of α_i where of course $s + t = m$. We replace the order $1, 2, \dots, m$ by the order $\sigma(1) = b_1, \sigma(2) = b_2, \dots, \sigma(s) = b_s, \sigma(s + 1) = c_1, \sigma(s + 2) = c_2, \dots, \sigma(m) = c_t$. We claim σ yields a total order for the new graph D . Consider a pair of rows $\sigma(p), \sigma(q)$ for $p < q$. If $1 \leq p < q \leq s$ or if $s + 1 \leq p < q \leq m$, then our contributing column has made no changes to the bucket $\{\sigma(p), \sigma(q)\}$ and so we have $\sigma(p) \rightarrow \sigma(q)$. If $1 \leq p \leq s$ and $s + 1 \leq q \leq m$, then either we begin with $\sigma(p) \rightarrow \sigma(q)$ and so our contributing column will not change the arc in D for $\{\sigma(p), \sigma(q)\}$ or we begin with $\sigma(q) \rightarrow \sigma(p)$ and so our contributing column will add to the bucket and the arc will be reversed in D . Thus we end with a new total order with the rows ordered using σ . By induction on the number of contributing columns processed, D is always a total order.

Let σ_i, σ_{i+1} be the total orders of D before and after processing α_i . We note that the number of contributions of α_i can be computed as follows. Let x be the smallest index so that entry $\sigma_i(x)$ of α_i is 1 and let y be the largest index so that entry $\sigma_i(y)$ of α_i is 0. Using that σ_i is a total order, we have that for $e < f$ that $\sigma_i(e) \rightarrow \sigma_i(f)$. For every z with $x < z < y$ we either have a 1 in α_i in row $\sigma_i(z)$ and so a contribution on the pair $\sigma_i(x), \sigma_i(z)$ or we have a 0 in α_i in row $\sigma_i(z)$ and so a contribution on the pair $\sigma_i(z), \sigma_i(y)$. Also we have a contribution from the pair of rows $\sigma_i(x), \sigma_i(y)$. Then the number of contributions from α_i is at least $y - x$.

The number of columns which would be filler columns with respect to σ_{i+1} is exactly $m + 1$ by Lemma 1.4 but we need only count those new filler columns that were not filler columns with respect to σ_i . It is easy to count that the number of possible filler columns with respect to σ_{i+1} that were not filler columns with respect to σ_i is at most $y - x$. Of course one of the new filler columns with respect to σ_{i+1} is in fact the contributing column α_i . Thus the arithmetic works out perfectly: the $y - x$ (or more) contributions from α_i yield $y - x$ credits to pay for α_i as well as the new filler columns with respect to σ_{i+1} . We will also need up to $m + 1$ credits to pay for the possible filler columns associated with the initial order (namely columns contained in $[T_m | \mathbf{0}_m]$). ■

4 Striped

Proof of Theorem 1.5: Let F be the matrix in (3). We may use Corollary 5.3 to establish the lower bound for $\text{fs}(m, F)$. For the upper bound, partition A into blocks $A = [A_0|A_1|A_2|\cdots]$, where each block A_i (except A_0) consists of a contributing column α_i followed by any number of filler columns. A_0 has no contributing columns and hence may have no columns. We have that A_i has neither submatrix $\begin{bmatrix} 1 & 0 \\ 1 & 0 \end{bmatrix}$ nor $\begin{bmatrix} 0 & 1 \\ 0 & 1 \end{bmatrix}$ and, as noted in describing the proof of Lemma 2.2 for $\|A_i\| \geq 3$, we can identify A_i as a sunflower or inverse sunflower and uniquely identify the centre X and the centre column $\psi(A_i)$ (which need not be present in A_i or A). For a sunflower, we define the *petals* $P(A_i) \subseteq [m] \setminus X$ of all rows for which there is a column of A_i with a 1 in the row (and all 1's on the rows of X). For an inverse sunflower replace 1 by 0 and continue to call them petals.

If $\|A_i\| = 2$ then either there is a single row where the two columns differ or there are two rows where the two columns differ. Identifying such a situation as a sunflower or inverse sunflower is not so helpful. Similarly for $\|A_i\| = 1$. Thus we define

$$P(A_i) = \emptyset \text{ if } \|A_i\| \leq 2. \quad (5)$$

We proceed with a slightly different form of amortized analysis. We will process the blocks A_i as whole units. In general we wish each column to be paid for with a credit but allow up to m columns to be unpaid for. We create a *debt structure* that keeps track of all columns we have not yet paid for (except for two special cases noted below when considering A_0 for which up to 2 columns may not be paid for). This debt structure consists of a set $U \subseteq [m]$ and a function ϕ from elements of U to columns of A so that for $u \in U$, column $\phi(u)$ has not yet been paid for. Thus, the size of our debt at any point will be bounded by m . As the debt structure changes, we show below that we are able to pay for all the deleted columns with credits. By this reasoning, we obtain the bound (4) where the m comes from the debt structure and the 2 comes from the special cases.

It suffices for each contributing column to assign 4 *special* credits plus a number of credits equal to twice the number of contributions. Note that we would not need special credits if we simply counted each contribution from a contributing column as 6 credits, and it is this overcount which yields the upper bound of (4). We are saying that the number of paid columns is at most 6 times the number of contributions and we have a bound on the number of contributions. The rest of the proof shows that the debt structure works as claimed.

We have a *split structure* for each A_i . For $\|A_i\| \geq 3$ we define $J(A_i)$ (respectively $Z(A_i)$) to be the rows which have a 1 (resp. 0) in the centre column $\psi(A_i)$. For $\|A_i\| \leq 2$, we select the first column α_i of A_i and similarly define $J(A_i)$ (respectively $Z(A_i)$) to be the rows which have a 1 (resp. 0) in the column α_i . In either case this yields a partition $J(A_i) \cup Z(A_i) = [m]$. The notations J, Z are in analogy with the use of J to represent the matrix of 1's and Z to represent the matrix of 0's. On every pair of rows u, v in $J(A_i)$, we are looking for $\begin{bmatrix} 0 \\ 0 \end{bmatrix}$ (the top of the stack for the bucket u, v is $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$). On every pair of rows in $Z(A_i)$, the reverse is true. Since the first column of F is $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$, we find it convenient to set $J(A_0) = \emptyset$ and $Z(A_0) = [m]$.

When processing A_0 , we set $U_0 = P(A_0)$ with $\phi(u)$ being the column associated with the petal in A_0 . But this may leave up to 2 columns of A_0 not in the debt structure. For $\|A_0\| \leq 2$ we have $U_0 = P(A_0) = \emptyset$ and so in this special case up to 2 columns are not in the debt structure and must be paid for in some other way. For $\|A_0\| \geq 3$, the centre column if present in A_0 is not added to the debt structure and so in this special case up to 1 column must be paid for in some other way. This gives the special cases noted above yielding the 2 in the bound (4).

We update the debt structure from U_{i-1}, ϕ_{i-1} to U_i, ϕ_i in a two stage process. First we delete some elements x of U_{i-1} and their corresponding columns $\phi_{i-1}(x)$. For each pair of rows u, v for which the contributing column α_i adds a contribution to the bucket $\{u, v\}$, we delete both u and v and the columns $\phi_{i-1}(u), \phi_{i-1}(v)$ (if present in the debt structure) from the debt structure. If $|J(A_{i-1}) \cap Z(A_i)| = 1$ and $\{w\} = J(A_{i-1}) \cap Z(A_i) \subseteq U_{i-1}$, then we delete w and $\phi_{i-1}(w)$ from the debt structure. If $|Z(A_{i-1}) \cap J(A_i)| = 1$ and $\{t\} = Z(A_{i-1}) \cap J(A_i) \subseteq U_{i-1}$, then we delete t and $\phi_{i-1}(t)$ from the debt structure. Then, for $\|A_i\| \geq 3$, we add $P(A_i)$ and all the petal columns from A_i to the debt structure.

After completing this process for A_i , then for $u \in U_i$,

$$\phi_i(u) = \begin{cases} \text{petal column of } A_i \text{ for row } u & \text{if } u \in P(A_i), \|A_i\| \geq 3, \\ \phi_{i-1}(u) & \text{otherwise.} \end{cases}$$

After processing block A_i , the debt structure consists of a set of rows $U_i \subseteq [m]$, together with a mapping ϕ_i from U_i to a petal column contained in $[A_1|A_2|\cdots|A_i]$. Recalling that A is simple yields that column $\phi_i(u)$ also identifies a unique block A_j for some $j \leq i$ which contains the column $\phi_i(u)$. Also, we only add columns to the debt structure when they are petal columns and so u is a petal in $P(A_j)$ and $\phi_i(u)$ corresponds to the column of A_j that yields the petal u . Each row $u \in U_i$ indicates that the column $\phi_i(u)$ has not yet been paid for by a credit. As we process blocks of A , we need that every column not yet paid for is remembered in our debt structure.

Note that we will need to immediately pay for the centre column $\psi(A_i)$, if it exists in A_i , since it does not correspond to a petal and so cannot enter the debt structure (there is no row to point to $\psi(A_i)$). Similarly we will have to immediately pay for all the columns of A_i in the case $\|A_i\| \leq 2$ since these columns do not enter the debt structure ($P(A_i) = \emptyset$).

We may use the two credits associated with a contribution of α_i to bucket $\{u, v\}$ to pay for the two columns $\phi_{i-1}(u), \phi_{i-1}(v)$ deleted from the debt structure. We use up to two special credits from α_i to either pay for the columns of A_i when $\|A_i\| \leq 2$ or to pay for the centre column (if present) for $\|A_i\| \geq 3$. We also use up to two special credits to pay for the deletion of $\phi_{i-1}(w)$ and $\phi_{i-1}(t)$ (if they are present) with w, t as described above. Thus if u moves from ‘ Z ’ to ‘ J ’ (namely $u \in Z(A_{i-1}) \cap J(A_i)$) then u and $\phi_{i-1}(u)$ are deleted from the debt structure in the first stage of the update (u may reenter U_i in the second stage if $u \in P(A_i)$). Similarly if u moves from ‘ J ’ to ‘ Z ’ (namely $u \in J(A_{i-1}) \cap Z(A_i)$) then u and $\phi_{i-1}(u)$ are deleted from the debt structure in the first stage of the update. In the worst case we need to assign to each contributing column 4 special credits plus 2 credits for each contribution.

The rest of our proof is to ensure that for rows $u \in P(A_i)$, when we add the petal column $\phi_i(u)$ from A_i to the debt structure, we have already deleted the column $\phi_{i-1}(u)$ for $u \in U_{i-1}$. Assume $u \in P(A_i)$ has column $\phi_{i-1}(u)$ still in the debt structure when we attempt to add the petal column of A_i associated to row u . Without loss of generality, we assume that A_i is a sunflower. Thus $u \in Z(A_i)$ (u will be a petal in the sunflower).

The column $\phi_{i-1}(u)$ belongs to A_j for some $j \leq i-1$. Let $\gamma = \phi_{i-1}(u) = \phi_j(u)$. If $u \in J(A_j) \cap Z(A_i)$, then at some intermediate stage we will have moved u from ‘ J ’ to ‘ Z ’ and γ will have been deleted at that point. Therefore we deduce that $u \in Z(A_j)$. Given that γ enters the debt structure while processing A_j we deduce that $\|A_j\| \geq 3$. Also since $u \in Z(A_j)$ and u must be a petal then A_j is a sunflower. If $J(A_j) = J(A_i)$, then γ is also a column of A_i (since $u \in P(A_i)$), contradicting that we have a simple matrix. Thus $J(A_j) \neq J(A_i)$.

Assume there is a row $v \in J(A_j) \setminus J(A_i) = J(A_j) \cap Z(A_i)$. Then we have $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ at the top of the stack in the bucket u, v after processing A_j (by the column γ in A_j) and we have $\begin{bmatrix} 0 \\ 0 \end{bmatrix}$ at the top of the stack in the bucket u, v after processing $A_1 A_2 \cdots A_i$ (since $u, v \in Z(A_i)$). Thus we will have a contribution on rows u, v by the time we have processed A_i and so we would have deleted u (and perhaps v if it is in the debt structure) and column γ from the debt structure U_{t-1} for some t with $j < t \leq i$. Thus we have deleted γ by the time we wish to add $P(A_i)$ to the debt structure.

Similarly assume there is a $v \in J(A_i) \setminus J(A_j) = J(A_i) \cap Z(A_j)$. Then we have $\begin{bmatrix} 0 \\ 0 \end{bmatrix}$ at the top of the stack in the bucket u, v after processing $A_1 A_2 \cdots A_j$ ($u, v \in Z(A_j)$) and we have $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ at the top of the stack in the bucket u, v after processing $A_1 A_2 \cdots A_i$ (by petal column for u in A_i). Again we will have deleted γ . This contradiction to our choice of u shows that the debt structure is always updated successfully. ■

We do not expect our coefficient of $6(\ell-1)\binom{m}{2}$ in Theorem 1.5 to be best possible and instead expect the lower bound to be closer to the truth. Computer searches suggest $(\ell-1)m^2/2$ will be the leading term. Our proof would not allow us to reduce the 6 to a number below 2.

5 Constructions

We offer two construction techniques (Proposition 5.1, Proposition 5.4) helpful in finding matrices avoiding a given submatrix.

Proposition 5.1 *Let F be a $k \times l$ $(0,1)$ -matrix with a non-constant bottom row. Let $F = [G|\phi|H]$, where ϕ is any column of F and the blocks G, H may be empty. Then*

$$fs(m, F) \geq fs(m-1, [G|\phi]) + fs(m-1, [\phi|H]). \quad (6)$$

Proof: Since F has a non-constant bottom row, either $[G|\phi]$ contains a 1 in its bottom row and $[\phi|H]$ contains a 0 in its bottom row, or vice-versa. Without loss of generality assume the former case. Find an $(m-1) \times fs(m-1, [G|\phi])$ simple matrix B with no

submatrix $[G|\phi]$ and an $(m-1) \times \text{fs}(m-1, [\phi|H])$ simple matrix C with no submatrix $[\phi|H]$ and consider

$$A = \left[\begin{array}{c|c} B & C \\ \hline 0 \cdots 0 & 1 \cdots 1 \end{array} \right].$$

Suppose A contains $F = [G|\phi|H]$ as a submatrix where the i th column of A contains the column of F corresponding to ϕ . Since F contains a $[10]$ in its bottom row, F must lie above the bottom row of A . If column i occurs in the first $\text{fs}(m-1, [G|\phi])$ columns of A then B contains $[G|\phi]$ as a submatrix, a contradiction. If column i occurs in the last $\text{fs}(m-1, [\phi|H])$ columns of A then C contains $[\phi|H]$ as a submatrix, a contradiction. Thus A cannot contain F as a submatrix, establishing (6). ■

We can generalize this as follows.

Corollary 5.2 *Let F be a $k \times \ell$ $(0,1)$ -matrix with a non-constant bottom row. Let $F = [\phi_0|F_1|\phi_1|F_2|\phi_2 \cdots \phi_{r-1}|F_r|\phi_r]$, where the ϕ_i 's are any columns of F and the F_i 's may be empty. Then*

$$\text{fs}(m, F) \geq \sum_{i=1}^r \text{fs}(m-r+1, [\phi_{i-1}|F_i|\phi_i]). \quad (7)$$

Proof: We use induction on r . When $r = 1$, the inequality is an equality. Now suppose the proposition holds with r replaced by $r-1$. Let $F = [\phi_0|F_1|\phi_1|F_2|\phi_2 \cdots \phi_{r-1}|F_r|\phi_r]$, where the ϕ_i 's are any columns of F and the F_i 's may be empty. Since F has a non-constant bottom row, either $[\phi_0 F_1 \phi_1 \cdots F_{r-1} \phi_{r-1}]$ or $[\phi_1 F_2 \cdots \phi_{r-1} F_r \phi_r]$ has a non-constant bottom row. Without loss of generality assume the former case. By Proposition 5.1 we have

$$\text{fs}(m, F) \geq \text{fs}(m-1, [\phi_0 F_1 \phi_1 \cdots F_{r-1} \phi_{r-1}]) + \text{fs}(m-1, [\phi_{r-1} F_r \phi_r]),$$

while by the induction hypothesis we have

$$\text{fs}(m-1, [\phi_0 F_1 \phi_1 \cdots F_{r-1} \phi_{r-1}]) \geq \sum_{i=1}^{r-1} \text{fs}(m-1-(r-1)+1, [\phi_{i-1} F_i \phi_i]).$$

Also $\text{fs}(m-1, [\phi_0 F_1 \phi_1 \cdots F_{r-1} \phi_{r-1}]) \geq \text{fs}(m-r+1, [\phi_0 F_1 \phi_1 \cdots F_{r-1} \phi_{r-1}])$ by Lemma 1.6. Thus (7) is true. ■

Our proof shows that $\text{fs}(m, F) \geq \sum_{i=1}^r \text{fs}(m-u_i, [\phi_{i-1}|F_i|\phi_i])$, where $(u_1, u_2, \dots, u_r)$ is some permutation of $(1, 2, 3, \dots, r-3, r-2, r-1, r-1)$ determined by our construction. We apply this construction idea to explore constructions for $2 \times \ell$ F . The case where F is either all 0's or all 1's can be found in [4]. We obtain good constructions for other 2-rowed F which may be close to optimal.

Corollary 5.3 *Let F be a $2 \times \ell$ $(0,1)$ -matrix. Assume F is not all 0's or all 1's. Let b be the number of pairs of consecutive columns of F which yield the submatrix $\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$ or $\begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}$. Then*

$$\text{fs}(m, F) \geq \binom{m}{2}(\ell-1-b) - \ell^2 m - 3\ell^2. \quad (8)$$

Proof: If F has all identical columns then the columns are either all $\begin{bmatrix} 1 \\ 0 \end{bmatrix}$ or all $\begin{bmatrix} 0 \\ 1 \end{bmatrix}$; in that case, $b = \ell - 1$ and so (8) is trivial. Otherwise, assume without loss of generality that F has a non-constant bottom row. Let $G_1, G_2, \dots, G_{\ell-1}$ be the pairs of consecutive columns of F . By Corollary 5.2 we have

$$\text{fs}(m, F) \geq \sum_{i=1}^{\ell-1} \text{fs}(m - \ell + 2, G_i).$$

By the results of Anstee and Füredi [4] and Frankl, Füredi, and Pach [5], for each G_i we have

$$\text{fs}(m - \ell + 2, G_i) \geq \binom{m - \ell + 2}{2} + (m - \ell + 2) + 1 \geq \binom{m - \ell}{2} + (m - \ell) + 1$$

unless G_i is $\begin{bmatrix} 1 & 1 \\ 0 & 0 \end{bmatrix}$ or $\begin{bmatrix} 0 & 0 \\ 1 & 1 \end{bmatrix}$, in which case we have instead

$$\text{fs}(m - \ell + 2, G_i) \geq 2(m - \ell + 2) \geq 2(m - \ell).$$

So

$$\text{fs}(m, F) \geq (\ell - 1 - b) \left(\binom{m - \ell}{2} + (m - \ell) + 1 \right) + 2b(m - \ell),$$

from which (8) follows. ■

This construction establishes the lower bounds of Theorems 1.4 and 1.5. However, for a matrix such as

$$F = \begin{bmatrix} 1 & 1 & 0 & 0 \\ 0 & 0 & 1 & 1 \end{bmatrix},$$

the construction yields the lower bound

$$\text{fs}(m, F) \geq \binom{m}{2} - O(m).$$

Our best upper bound is

$$\text{fs}(m, F) \leq 5 \binom{m}{2} + m + 1,$$

obtained from using the 2×6 forbidden submatrix $\begin{bmatrix} 101010 \\ 010101 \end{bmatrix}$. We expect the smaller bound to be correct and computer searches have supported this.

In order to apply Corollary 5.2 to more general F , it would be helpful to have good constructions for any $k \times 2$ matrix (obtained as a pair of consecutive columns of F). We already have bounds for each $k \times 2$ F consisting of a pair of identical columns [4]. The following handles $k \times 2$ F with two non-identical columns. We use A^T to refer to the transpose of A .

Proposition 5.4 *Let F be a $k \times 2$ $(0,1)$ -matrix with two different columns. Then*

$$\text{fs}(m, F) \geq \binom{m}{k} + \binom{m}{k-1} + \dots + \binom{m}{0}. \quad (9)$$

Proof: Let

$$F = \begin{bmatrix} a_1 & * \\ \vdots & \vdots \\ a_p & * \\ 1 & 0 \\ * & b_1 \\ \vdots & \vdots \\ * & b_q \end{bmatrix} \quad (10)$$

with $p + q + 1 = k$; either p or q may be 0. Let

$$\gamma = [a_1, a_2, \dots, a_p, 1, 0, b_1, b_2, \dots, b_q]^T.$$

Following Lemma 1.3, let A be the $m \times ((\binom{m}{k} + \binom{m}{k-1} + \dots + \binom{m}{0}))$ of all columns not containing γ as a submatrix. Then, as described in Lemma 1.3, every column α of A is an initial segment (from the top) of some instance of

$$\left((\bar{a}_1)^* a_1 (\bar{a}_2)^* a_2 \cdots (\bar{a}_p)^* a_p 0 \underset{\uparrow}{1} 1^* 0 (\bar{b}_1)^* b_1 (\bar{b}_2)^* b_2 \cdots (\bar{b}_q)^* \right)^T.$$

Let $g(\alpha)$ denote the index of the row in which the 1 indicated above occurs in α , or ∞ if the indicated 1 does not occur in α . Let B be the matrix obtained from A by reordering the columns of A by nonincreasing g , so that if column α occurs to the left of column β in B then $g(\alpha) \geq g(\beta)$. Now B achieves the bound (9). To show that B has no submatrix F , assume that F occurs in two columns α, β of B , α to the left of β and so $g(\alpha) \geq g(\beta)$. By definition of g , the 1 in row $p+1$ of F occurs in or below row $g(\alpha)$ of α . On the other hand, the 0 in row $p+1$ of F occurs above row $g(\beta)$ of β , since otherwise β contains a copy of γ as a submatrix. Thus $g(\alpha) < g(\beta)$, a contradiction. ■

The construction used in the proof of Theorem 2.5 in [4] occurs as a special case when $a_1 = a_2 = \dots = a_p = 1$ and in addition the starred entries to their right (in (10)) are also 1.

References

- [1] R.P. Anstee, On a Conjecture concerning Forbidden Submatrices, *J. Combin. Math. and Combin. Comp.* **32**(2000), 185-192.
- [2] R.P. Anstee, A. Sali, Small Forbidden Configurations IV, *Combinatorica* **25**(2005), 503-518.
- [3] R.P. Anstee, M. Farber, Characterizations of Totally Balanced Matrices, *J. Algorithms.* **5**(1984), 215-230.
- [4] R.P. Anstee, Z. Füredi, Forbidden Submatrices, *Discrete Math.* **62** (1986), 225-243.
- [5] P. Frankl, Z. Füredi and J. Pach, Bounding one-way differences, *Graphs and Combinatorics*, **3** (1987), 341-347.

- [6] A.J. Hoffman, A.W.J. Kolen, and M. Sakarovitch, Totally-balanced and greedy matrices, *SIAM J. Algebraic Discrete Methods* **7**(1986), 348-357.
- [7] G. Tardos, On 0-1 matrices and small excluded submatrices, *J. Combin. Th. Ser. A* **111** (2005), 266-288.