Forbidden configurations: Finding the number predicted by the Anstee-Sali conjecture is NP-hard

Miguel Raggi *

*Centro de Ciencias Matemáticas and Escuela Nacional de Estudios Superiores
Universidad Nacional Autónoma de México*

Received 31 January 2013, accepted 23 June 2014, published online 24 January 2015

Abstract

Let F be a (possibly non-simple) hypergraph and let $\text{forb}(m, F)$ denote the maximum number of edges a simple hypergraph with m vertices can have if it doesn't contain F as a subhypergraph. A conjecture of Anstee and Sali predicts the asymptotic behaviour of $\text{forb}(m, F)$ for fixed F . In this paper we prove that even finding this predicted asymptotic behaviour is an NP-hard problem, meaning that if the Anstee-Sali conjecture were true, finding the asymptotics of $\text{forb}(m, F)$ would be NP-hard.

Keywords: Forbidden configuration, hypergraph, trace, NP-hard, NP-complete, Anstee-Sali conjecture.

Math. Subj. Class.: 05D05, 68R01

1 Introduction

This paper considers an extremal problem in hypergraph theory that results as a natural generalization of Turán's famous problem.

Some of the most celebrated extremal results are those of Erdős, Stone and Simonovits ([11], [10]). They consider the following problem: Given $m \in \mathbb{N}$ and a graph F , find the maximum number of edges in a graph on m vertices that avoids having a subgraph isomorphic to F .

There are a number of ways to generalize this to hypergraphs. A k -uniform hypergraph is one in which each edge has size k . Some view k -uniform hypergraphs as the most natural generalization of a graph (a graph is a 2-uniform hypergraph) and one might also generalize

*Research supported by DGAPA Postdoctoral scholarship, UNAM.

E-mail address: mraggi@gmail.com (Miguel Raggi)

the forbidden subgraph problem to a forbidden k -uniform subhypergraph problem. There are both asymptotic and exact results (e.g. [9], [13], [12]).

Forbidden Configurations is a different (but also natural) generalization that is studied mainly by Richard Anstee and his collaborators. We consider the following problem: Given $m \in \mathbb{N}$ and a hypergraph F , find the maximum number of edges in a *simple* hypergraph (i.e. no repeated edges) H on m vertices that avoids having a subhypergraph isomorphic to F . Surveys on the topic can be found in [1] and [14].

We find it convenient to use the language of matrices to describe hypergraphs: Each column of a $\{0, 1\}$ -matrix can be viewed as an incidence vector on the set of rows.

Definition 1.1. If α is a column and A a matrix, define $\lambda(\alpha, A)$ to be the multiplicity of α in A . Define a matrix to be **simple** if it is a $\{0, 1\}$ -matrix with no repeated columns (that is, if $\lambda(\alpha, A) \leq 1$ for every column α).

Note that an $m \times n$ simple matrix corresponds to a **simple hypergraph** (or **set system**) on m vertices with n distinct edges, where we allow the “empty edge”.

Definition 1.2. When A is a $\{0, 1\}$ -matrix, we denote by $\|A\|$ the number of columns in A (which is the cardinality of the associated set system).

Definition 1.3. Let A and B be $\{0, 1\}$ -matrices with the same number of rows. Define the **concatenation** $[A|B]$ to be the configuration that results from taking all columns of A together with all columns of B . For $t \in \mathbb{N}$, we define the product

$$t \cdot A := \underbrace{[A \mid A \mid \cdots \mid A]}_{t \text{ times}}.$$

Our objects of study are $\{0, 1\}$ -matrices with row and column order information stripped from them.

Definition 1.4. Two $\{0, 1\}$ -matrices are said to be equivalent if one is a row and column permutation of the other. An equivalence class is called a **configuration**.

Abusing notation, we will commonly use matrices and their corresponding configurations interchangeably.

Definition 1.5. For a configuration F and a $\{0, 1\}$ -matrix A (or a configuration A), we say that F is a **subconfiguration** of A , and write $F \prec A$ if there is a representative of F which is a submatrix of A . We say A **has no configuration** F (or **doesn't contain F as a configuration**) if F is not a subconfiguration of A . Let $\text{Avoid}(m, F)$ denote the set of all simple matrices on m -rows with no subconfiguration F .

Our main extremal problem is to compute

$$\text{forb}(m, F) = \max_A \{\|A\| : A \in \text{Avoid}(m, F)\}.$$

Perhaps some examples are useful:

- $\text{forb}\left(m, \begin{bmatrix} 1 \\ 1 \end{bmatrix}\right) = m + 1$, since we can take all columns with at most one 1.

- $\text{forb}\left(m, \begin{bmatrix} 1 \\ 0 \end{bmatrix}\right) = 2$, since we may only take the column of 1's and the column of 0's (i.e. the empty set and the complete set).
 - $\text{forb}\left(m, \begin{bmatrix} 1 & 1 \\ 1 & 1 \end{bmatrix}\right) = \binom{m}{2} + \binom{m}{1} + \binom{m}{0}$, by taking all columns with at most two 1's.
- The proof that this is indeed the maximum is easy and can be found in [14].

Let A^c denote the $\{0, 1\}$ -complement of A (replace every 0 in A by a 1 and every 1 by a 0). Note that $\text{forb}(m, F) = \text{forb}(m, F^c)$.

Remark 1.6. Let F and G be configurations such that $F \prec G$. Then $\text{forb}(m, F) \leq \text{forb}(m, G)$.

We say a column α has **column sum** t if it has exactly t ones. Let $\mathbf{0}_m$ denote the column with m rows, all of them zeros. Similarly, let $\mathbf{1}_m$ denote the column of m ones.

For a set of rows S , we let $A|_S$ denote the submatrix of A given by restricting the rows of A to only those in S .

An important general result due to Füredi applies to simple and to non-simple configurations.

Theorem 1.7 (Z. Füredi). *Let F be a given k -rowed $\{0, 1\}$ -matrix. Then $\text{forb}(m, F)$ is in $O(m^k)$.*

We desire more accurate asymptotic bounds. Anstee and Sali conjectured that the best asymptotic bounds can be achieved with certain *product constructions*.

Definition 1.8. Let A and B be $\{0, 1\}$ -matrices. We define the **product** $A \times B$ by taking each column of A and putting it on top of every column of B . Here is an example of a product:

$$A = \begin{bmatrix} 0 & 1 & 1 \\ 0 & 0 & 1 \end{bmatrix}, B = \begin{bmatrix} 1 & 0 \\ 0 & 1 \end{bmatrix} \implies \begin{matrix} A \\ \times \\ B \end{matrix} = \left[\begin{array}{cc|cc|cc} 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 1 & 0 & 1 & 0 & 1 & 0 \\ 0 & 1 & 0 & 1 & 0 & 1 \end{array} \right].$$

Note that this is a well defined operation in configurations.

We are interested in asymptotic bounds for $\text{forb}(m, F)$. Let $\mathcal{I}_m$ be the $m \times m$ identity matrix, $\mathcal{I}_m^c$ be the $\{0, 1\}$ -complement of $\mathcal{I}_m$ (all ones except for the diagonal) and let $\mathcal{T}_m$ be the tower matrix: a matrix corresponding to a maximum chain in the partially ordered set of the power set of the vertices. For example,

$$\mathcal{T}_4 = \begin{bmatrix} 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 \end{bmatrix}.$$

Anstee and Sali conjectured that the asymptotically “best” constructions avoiding a single configuration would be products of $\mathcal{I}, \mathcal{I}^c$ and $\mathcal{T}$.

Definition 1.9. Let F be a configuration. Let

$$P_r(a, b, c) := \underbrace{\mathcal{I}_r \times \dots \times \mathcal{I}_r}_{a \text{ times}} \times \underbrace{\mathcal{I}_r^c \times \dots \times \mathcal{I}_r^c}_{b \text{ times}} \times \underbrace{\mathcal{T}_r \times \dots \times \mathcal{T}_r}_{c \text{ times}},$$

Define $X(F)$ to be the largest number such that there exist numbers $a, b, c \in \mathbb{N}$ with $a + b + c = X(F)$ such that for all $r \in \mathbb{N}$,

$$F \not\leq P_r(a, b, c).$$

Conjecture 1.10. [8] Let F be a configuration. Then $\text{forb}(m, F)$ is in $\Theta(m^{X(F)})$.

Observe that $X(F)$ is always an integer. Also note that $\|P_r(a, b, c)\| = r^{a+b} \cdot (r+1)^c \in \Theta(r^{X(F)})$, so by taking $r = \lceil m/X(F) \rceil$ (and perhaps deleting a constant number of rows and therefore columns in case $X(F) \nmid m$), we have that $\|P_r(a, b, c)\| \in \Omega(m^{X(F)})$. So the fact that $\text{forb}(m, F) \in \Omega(m^{X(F)})$ is built into the conjecture.

Thus, in order to prove the conjecture, all that would be required would be to prove that $\text{forb}(m, F) \in O(m^{X(F)})$ for every F . A disproof could be potentially easier, as only a counterexample would be required.

The conjecture has been proven for all $k \times \ell$ configurations F with $k \in \{1, 2, 3\}$ and many others cases in various papers. The proofs for $k = 2$ are in [4], for $k = 3$ in [4], [2], [8]. For $\ell = 2$, the conjecture was verified in [5]. For $k = 4$, all cases either when the conjecture predicts a cubic bound for F or when F is simple were completed in [3]. For $k = 4$ and F non-simple, there are three boundary cases with quadratic bounds, one of which is established in [6]. For $k \in \{5, 6\}$ some results can be found in [7].

Anstee has long conjectured that even finding $X(F)$ given F was not a trivial task, and the specific question of its NP-hardness was conjectured in [1], [14]. In this paper we settle this question: finding X is indeed NP-hard. We also note that one of the decision versions associated with this optimization problem is NP-complete, adding this function to the long list of functions known to be NP-complete, with the interesting plus that this function is conjectured to give the *exponent* of the asymptotic growth of forb .

For relatively small configurations F we have a computer program that yields the answer (relatively) quickly. The source code (in C++) can be freely downloaded from:

<http://matmor.unam.mx/~mraggi/>

This program can compute $X(F)$ for F having less than ~ 10 rows in just a few minutes. This task takes merely exponential time, not doubly exponential (as it is often the case with forbidden configuration problems). This program was written to perform many tasks other than finding $X(F)$. A description of the algorithms used can be found in [14].

2 Results

There are two natural decision problems associated with $X(F)$: Given F and k as inputs,

1. Is it true that $X(F) < k$?
2. Is it true that $X(F) \geq k$?

We prove that the first of the two decision problems is in NP by exhibiting a certificate which can be checked in polynomial time.

The main result of this paper is the following:

Theorem 2.1. *Finding $X(F)$ is an NP-hard problem. In other words, should a polynomial-time algorithm exist for finding $X(F)$ given F , then every problem in NP could be solved in polynomial time. Furthermore, the problem “given F and k , is $X(F) < k$?” is in NP.*

Before proving this theorem, we need a few results.

Proposition 2.2. *Let F be a configuration with n rows. Then $X(F) \leq n$.*

Proof. Indeed, assume for the sake of contradiction that a, b , and c are such that $a + b + c = n + 1$ and $F \not\prec P_r(a, b, c)$. Every configuration of $P_r(a, b, c)$ contains $[01]$, therefore any $\{0, 1\}$ -column can be formed with the first n configurations of $P_r(a, b, c)$. The $n + 1$ -th matrix ensures the columns of F with high multiplicity get repeated as many times as needed. $\square$

This observation in particular implies that if a polynomial time algorithm existed for any of the two decision versions of the problem, then we’d have a polynomial time algorithm for finding $X(F)$, which together with Theorem 2.1 would make the decision version of finding $X(F)$ an NP-complete problem.

A simple (but surprising) corollary of Conjecture 1.10, if it were true, would be that repeating columns more than twice in F has no effect on the asymptotic behavior of $\text{forb}(m, F)$. In other words, assuming the conjecture were true, the multiplicity of a column in a configuration would not affect the asymptotic bound and, asymptotically, it would only matter if a column is not there (has multiplicity 0), appears once (has multiplicity 1), or appears “multiple times” (has multiplicity 2 or more). Formally,

Proposition 2.3. *Let $F_t = [G|t \cdot H]$ with G and H simple $\{0, 1\}$ -matrices that have no columns in common. Then $X(F_2) = X(F_t)$ for all $t \geq 2$. In particular, if the conjecture were true, then $\text{forb}(m, F_t)$ and $\text{forb}(m, F_2)$ would have the same asymptotic behavior.*

Proof. It suffices to show that given t, G, H, a, b and c , there exists an R such that for every $r \geq R$, we have

$$F_2 = [G|2 \cdot H] \prec P_r(a, b, c) \iff F_t = [G|t \cdot H] \prec P_r(a, b, c).$$

Since $F_2 \prec F_t$, we only need to prove that if $F_2 \prec P_r(a, b, c)$ for some r , then $F_t \prec P_R(a, b, c)$ for some R . Suppose then F_2 is contained in the product $P_r(a, b, c)$ for some r . The idea is to find a subconfiguration of $P_r(a, b, c)$ in which there are some columns with multiplicity 1, and for the columns with multiplicity 2 or more, the multiplicity depends on r , and goes to infinity as r goes to infinity. Then we may take r to be large enough so that the multiplicity of any one column (with multiplicity of 2 or more) is larger than t .

Let x be the number of rows of F_t . Notice the following three facts, which include definitions for $E_{\mathcal{I}}$, $E_{\mathcal{I}^c}$ and $E_{\mathcal{T}}$.

$$\begin{aligned} E_{\mathcal{I}}(x, r) &:= [(r - x) \cdot \mathbf{0}_x \mid \mathcal{I}_x] \prec \mathcal{I}_r \\ E_{\mathcal{I}^c}(x, r) &:= [(r - x) \cdot \mathbf{1}_x \mid \mathcal{I}_x^c] \prec \mathcal{I}_r^c \\ E_{\mathcal{T}}(x, r) &:= \left\lfloor \frac{r}{x} \right\rfloor \cdot \mathcal{T}_x \prec \mathcal{T}_r. \end{aligned}$$

The first and second facts are easy to see; just take any subset of x rows from $\mathcal{I}_r$ or $\mathcal{I}_r^c$. The third statement is true by taking the $\lfloor r/x \rfloor$ -th row of $\mathcal{T}_r$, the $2\lfloor r/x \rfloor$ -th row of $\mathcal{T}_r$, etc.,

up to the $x \lfloor r/x \rfloor$ -th row. For example, if $r = 5$ and $x = 2$, we may take the second and fourth row from $\mathcal{T}_5$:

$$\mathcal{T}_5 = \begin{bmatrix} 0 & 1 & 1 & 1 & 1 & 1 \\ 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \implies \mathcal{T}_5|_{\{2,4\}} = \begin{bmatrix} 0 & 0 & 1 & 1 & 1 & 1 \\ 0 & 0 & 0 & 0 & 1 & 1 \end{bmatrix} = E_{\mathcal{T}}(2, 5)$$

Note that in the three configurations $E_{\mathcal{I}}(x, r)$, $E_{\mathcal{I}^c}(x, r)$ and $E_{\mathcal{T}}(x, r)$, we have that there are some columns of multiplicity 1 and there are some columns for which their multiplicity can be made as large as we wish by making r large. Formally, let $E(x, r)$ be one of $E_{\mathcal{I}}(x, r)$ or $E_{\mathcal{I}^c}(x, r)$ or $E_{\mathcal{T}}(x, r)$. We have that for every x -rowed column α there are three possibilities: either $\lambda(\alpha, E(x, r)) = 0$ for all r , or $\lambda(\alpha, E(x, r)) = 1$ for all r , or $\lim_{r \rightarrow \infty} \lambda(\alpha, E(x, r)) = \infty$.

If α is a column for which $\lim_{r \rightarrow \infty} \lambda(\alpha, E(x, r)) = \infty$, we may conclude that there is an R for which $\lambda(\alpha, E(x, r)) \geq t$ for every $r \geq R$.

Since F_2 is contained in $P_r(a, b, c)$ for some r , the columns in H will have multiplicity at least 2 in some subset of the rows of $P_r(a, b, c)$. We see that F_t is also a subconfiguration of $P_R(a, b, c)$. $\square$

3 Proof of the Main Theorem

We now prove the main theorem 2.1.

Proof. First we prove that the decision problem has a certificate which can be checked in polynomial time. A certificate that indeed $X(F) < k$ would have to be a proof that $F \prec P_r(a, b, c)$ for each triple a, b, c for which $a + b + c = k$. Note that there are at most a quadratic (with respect to the number of rows) number of a, b, c 's which satisfy the equation, since the question has a trivial “yes” answer when k is more than the number of rows (Proposition 2.2).

Given F and A configurations, one can easily construct a certificate that a configuration F is indeed a subconfiguration of a configuration A : explicitly state which permutation of F appears in exactly which rows and columns of A . For the case $A = P_r(a, b, c)$, a certificate only needs to specify which rows of F go inside which factors, so at most a quadratic number of these certificates-of-being-a-subconfiguration suffice.

We now prove that finding $X(F)$ is NP-hard. Suppose there existed some polynomial-time algorithm that finds $X(F)$ given F . We shall prove that there would then exist a polynomial time algorithm for GRAPH COLORING. Suppose we are given a graph G and we wish to find the minimum number of colors for which there exists a good coloring of the graph. We may assume that no isolated vertices exist.

The idea is to construct a 3-part matrix $F(G)$ in which the first two parts ensure there is no $\mathcal{T}$ or $\mathcal{I}^c$ in a maximum product of the form $P_r(a, b, c)$ with no subconfiguration $F(G)$, and the last part is constructed so that a partition into $\mathcal{I}'$ s produces a partition of the vertices of the graph into independent sets and vice-versa.

Suppose G has n vertices and e edges. Let M be a large number with $M \geq n + 2$ and let S be the incidence matrix of G (i.e., the edges of G are encoded as columns with two

1's corresponding to the vertices that belong to the edge). Construct the following simple matrix:

$$F(G) := \left[\begin{array}{c|c} 1 & \mathbf{1}_M \mathcal{I}_M^c \\ \hline 1 & \mathcal{T}_M \\ \hline S & 0 \end{array} \right].$$

Clearly we can construct $F(G)$ in polynomial time (with respect to the number of vertices of G). We prove now that we have $\chi(G) = X(F) - 2M + 1$, which in turn would yield a polynomial time algorithm for GRAPH COLORING, provided we had a polynomial time algorithm for finding $X(F)$.

Now, let us study the possibilities for a product of type $P_r(a, b, c)$ that does not have $F(G)$ as a subconfiguration for any r . If $b \neq 0$, then we could place all of $[1|\mathcal{I}_M^c]$ in the $\mathcal{I}^c$ part of $P_r(a, b, c)$, so $a + b + c$ would be at most $1 + M + n$ (using Proposition 2.2). The same conclusion holds when $c \neq 0$. But if we let $b = c = 0$, $P_r(a, 0, 0)$ is just a product of $\mathcal{I}$'s, so let us calculate how many $\mathcal{I}$'s we can multiply together and still not create a subconfiguration $F(G)$.

In order for $F(G)$ to be a part of a product of $\mathcal{I}$'s, every row of $[1|\mathcal{I}_M^c]$ and $[1|\mathcal{T}_M]$ must be in a separate factor $\mathcal{I}$, since there is no $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ in $\mathcal{I}$ (and also separate from the rows of $[S|0]$, since we are assuming G has no isolated vertices).

Then two rows of the $[S|0]$ part can be in the same $\mathcal{I}$ if and only if there is no $\begin{bmatrix} 1 \\ 1 \end{bmatrix}$ in those two rows, which, in terms of the graph, means there is no edge between those two vertices. On the other hand, a product of $\chi(G)$ identity matrices contains the incidence matrix of a complete $\chi(G)$ -multipartite graph. In other words, partitioning $[S|0]$ into $\mathcal{I}$'s is equivalent to partitioning the vertices of G into independent sets. So if the graph G cannot be colored with $\chi(G) - 1$ colors and this is the maximum, this means that $X(F) = a = 2M + \chi(G) - 1 \geq n + M + 1$. Then $\chi(G) = X(F(G)) - 2M + 1$. $\square$

Acknowledgements The author would like to thank the anonymous referees for their invaluable comments that doubtlessly improved the presentation of the paper.

References

- [1] R. Anstee, A survey of forbidden configuration results, <http://www.combinatorics.org/ojs/index.php/eljc/article/view/DS20>.
- [2] R. Anstee, R. Ferguson and A. Sali, Small forbidden configurations II, *Electronic Journal of Combinatorics* **8** (2001), R4 25pp.
- [3] R. Anstee and B. Fleming, Two refinements of the bound of Sauer, Perles and Shelah and Vapnik and Chervonenkis, *Discrete Mathematics* **310** (2010), 3318–3323.
- [4] R. Anstee, J. Griggs and A. Sali, Small forbidden configurations, *Graphs and Combinatorics* **13** (1997), 97–118.
- [5] R. Anstee and P. Keevash, Pairwise intersections and forbidden configurations, *European Journal of Combinatorics* **27** (2006), 1235–1248.
- [6] R. Anstee, M. Raggi and A. Sali, Evidence for a forbidden configuration conjecture: One more case solved, *Discrete Mathematics* **312** (2012), 2720 – 2729, doi:10.1016/j.disc.2012.02.020.
- [7] R. Anstee, M. Raggi and A. Sali, Forbidden configurations: Boundary cases, *European Journal of Combinatorics* (2014), 51.

- [8] R. Anstee and A. Sali, Small forbidden configurations IV, *Combinatorica* **25** (2005), 503–518.
- [9] D. de Caen and Z. Füredi, The maximum size of 3-uniform hypergraphs not containing a Fano plane, *Journal of Combinatorial Theory, Series B* **78** (2000), 274–276.
- [10] P. Erdős and M. Simonovits, A limit theorem in graph theory., *Studia Scientiarum Mathematicarum Hungarica* **1** (1966), 51–57.
- [11] P. Erdős and A. Stone, On the structure of linear graphs, *Bulletin of the American Mathematical Society* **52** (1946), 1089–1091.
- [12] Z. Füredi, Turán type problems, *Surveys in Combinatorics (Proc. of the 13th British Combinatorial Conference)*, ed. A.D. Keedwell, Cambridge Univ. Press. London Math. Soc. Lecture Note Series **166** (1991), 253–300.
- [13] O. Pikhurko, An exact bound for Turán result for the generalized triangle, *Combinatorica* **28** (2008), 187–208.
- [14] M. Raggi, *Forbidden Configurations*, Ph.D. thesis, University of British Columbia, 2011.