




ELSEVIER

Contents lists available at ScienceDirect

Linear Algebra and its Applications

www.elsevier.com/locate/laa


Normal limiting distributions for systems of linear equations in random sets

Juanjo Rué^{a,*}, Maximilian Wötzel^b

^a *Departament de Matemàtiques and Institut de Matemàtiques de la UPC (IMTech), Universitat Politècnica de Catalunya & Centre de Recerca Matemàtica & Barcelona Graduate School of Mathematics (BGSMath), Barcelona, Spain*

^b *Department of Mathematics, Radboud University, Nijmegen, the Netherlands*

ARTICLE INFO

Article history:

Received 3 December 2021

Accepted 6 May 2022

Available online 11 May 2022

Submitted by Z. Bao

MSC:

11D45

15A06

05B20

Keywords:

Systems of linear equations

Central limit theorem

Method of moments

ABSTRACT

We consider the binomial random set model $[n]_p$ where each element in $\{1, \dots, n\}$ is chosen independently with probability $p := p(n)$. We show that for essentially all regimes of p and very general conditions for a matrix A and a column vector \mathbf{b} , the count of specific integer solutions to the system of linear equations $A\mathbf{x} = \mathbf{b}$ with the entries of \mathbf{x} in $[n]_p$ follows a (conveniently rescaled) normal limiting distribution. This applies among others to the number of solutions with every variable having a different value, as well as to a broader class of so-called non-trivial solutions in homogeneous strictly balanced systems. Our proof relies on the delicate linear algebraic study both of the subadjacent matrices and the corresponding ranks of certain submatrices, together with the application of the method of moments in probability theory.

© 2022 The Author(s). Published by Elsevier Inc. This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

* Corresponding author.

E-mail addresses: juan.jose.rue@upc.edu (J. Rué), maximilian.woetzel@ru.nl (M. Wötzel).

1. Introduction and main results

The study of the existence of solutions to linear equations in subsets of the integers (and more generally in additive or even non-abelian groups) has been a prominent topic not only in number theory, but also in extremal combinatorics, ergodic theory, functional analysis and theoretical computers science, among other research areas. A prototypical example of such investigations is Roth's Theorem [1], which proves (by using Fourier analytic means) that dense sets of integers always contain arithmetic progressions of length 3. This result was fully generalized by Szemerédi [2], who obtained a similar statement for arithmetic progressions of all (fixed) lengths.

Following this line of research, Frankl, Graham and Rödl [3, Theorem 2] proved similar results for homogeneous systems of linear equations. In particular, they proved that when A is a *density regular* matrix with integer entries (that is, the columns of A sum up to the zero vector), then any dense integer set T will contain solutions to the linear system $Ax = 0$. This result generalizes Szemerédi's Theorem, as arithmetic progressions of length k (or k -APs for short) can be encoded as solutions to the linear system $x_2 - x_1 = x_3 - x_2 = \dots = x_k - x_{k-1}$. Other classical equations, such as the Schur equation for sum-free sets (sets without solutions to the equation $x + y = z$), and Sidon sets (sets without non-trivial solutions to the equation $x + y = z + t$) also fit into this framework.

A very recent trend of investigation is to transfer these extremal results to a sparse setting, and hence to prove analogues on subsets which are dense relative to certain ambient sets. Specific examples for such sparse ambient sets are for instance the primes or perfect powers, but also the binomial random set $[n]_p$ obtained by choosing to include independently each element in $[n] = \{1, 2, \dots, n\}$ with probability p . Returning to the example of k -APs, define a set T to be (δ, k) -Szemerédi if every subset $U \subseteq T$ with at least $\delta|T|$ elements contains a k -AP. Thus, Szemerédi's Theorem shows that $[n]$ (in fact, every set T that is itself dense in $[n]$) is (δ, k) -Szemerédi for every δ and k . For the specific case of $k = 3$, Kohayakawa, Łuczak and Rödl [4] proved that for all $\delta > 0$ there exist constants $c, C > 0$ such that if $p \leq cn^{-1/2}$, then as n tends to infinity, $\mathbf{P}([n]_p \text{ is } (\delta, 3)\text{-Szemerédi})$ tends to 0, while if $p \geq Cn^{-1/2}$, it tends to 1. The existence of this threshold was extended to all values of k in [5,6], and later rediscovered in [7, 8] in the context of independent sets in hypergraphs. These ideas, named *hypergraph container method* were then used by Spiegel [9] and independently by Hancock, Staden and Treglown [10] to extend [3, Theorem 2] to the broadest class of linear systems possible. Similar techniques were used by Rué, Serra and Vena [11] to study random sparse analogues of [3, Theorem 2] in finite fields and more general configurations than linear systems of equations.

The main contribution of our paper follows this trend of research, with the aim of studying the *typical* number of solutions to linear systems of equations in random sets of integers. Our main theorem, Theorem 18 is technical and in order to formally state it we need to introduce a wide variety of algebraic definitions and lemmas, which will be the core of Section 2. However, we will now introduce some main definitions in order

to be able to formulate two important consequences of Theorem 18, namely Theorems 4 and 7.

Proper solutions and their distribution In order to properly state our results we need to introduce some notation. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ an integer matrix, and $\mathbf{b} \in \mathbb{Z}^r$ an integer vector. We write the *rank* of a matrix A by $\text{rk}(A)$. Define $S(A, \mathbf{b}) = \{\mathbf{x} \in \mathbb{Z}^m : A\mathbf{x} = \mathbf{b}\}$ as the set of integer solutions to the system of linear equations $A\mathbf{x} = \mathbf{b}$. Furthermore, if $Q \subset [m]$, let $A^Q \in \mathbb{Z}^{r \times |Q|}$ denote the $r \times |Q|$ matrix obtained by only keeping the columns of A indexed by Q . We notice that the rank of the empty matrix A^\emptyset is zero. We will identify tuples $\mathbf{x} = (x_1, \dots, x_m) \in \mathbb{Z}^m$ with the corresponding column vectors but abuse notation slightly by letting \mathbf{x}^Q denote the vector obtained by only keeping the *rows* of \mathbf{x} . This should not be confused with the notation \mathbf{x}^k , where k is a positive integer, which we will use to denote the vector

$$\mathbf{x}^k = \underbrace{(\mathbf{x}, \dots, \mathbf{x})}_{k \text{ times}} = (x_1, \dots, x_m, x_1, \dots, x_m, \dots, x_1, \dots, x_m) \in \mathbb{Z}^{km}.$$

Finally, if $\mathbf{x} = (x_1, \dots, x_m) \in \mathbb{Z}^m$ is an m -tuple, we will write $\{\mathbf{x}\}$ as a shorthand for the set $\{x_1, \dots, x_m\}$. In particular, if some entry in \mathbf{x} is repeated, then the cardinality of $\{\mathbf{x}\}$ is strictly less than m .

Roughly speaking, our main goal in this paper is to estimate $|S(A, \mathbf{b}) \cap [n]_p^m|$, that is, the number of solutions whose entries lie in $[n]_p$ for different regimes of p . In order for this to be a well-posed problem, we need some definitions from [12] and [13].

Definition 1. An integer matrix $A \in \mathbb{Z}^{r \times m}$ is said to be

- i) *positive* if there exists an integer solution to $A\mathbf{x} = \mathbf{0}$ having all positive entries:

$$S(A, \mathbf{0}) \cap \mathbb{N}^m \neq \emptyset,$$

and if for any pair of indices $i, j \in [m]$, $i \neq j$, there exists a solution $(x_1, \dots, x_m) \in S(A, \mathbf{0})$ satisfying $x_i \neq x_j$.

- ii) *abundant* if $\text{rk}(A) > 0$ and the removal of at most two columns does not change the rank of A :

$$\text{rk}(A^Q) = \text{rk}(A)$$

for any $Q \subset [m]$ satisfying $|Q| \geq m - 2$.

Sometimes the second requirement for positivity in Definition 1 is called *irredundancy*. For our applications, we will never consider positivity and irredundancy separately, and hence we have combined those two properties into a single one for expediency’s sake.

In general, considering any possible solution in $S(A, \mathbf{b})$ is a bit too lenient. For instance, in the 3-AP situation (whose associated matrix is $A = \begin{pmatrix} 1 & -2 & 1 \end{pmatrix}$), $S(A, \mathbf{0})$ also contains all the tuples (a, a, a) with $a \in \mathbb{Z}$, and those are clearly not of interest. To remedy this issue we need to introduce the easy notion of *proper* solutions to $A\mathbf{x} = \mathbf{b}$.

Definition 2. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, and $\mathbf{b} \in \mathbb{Z}^r$. Then the set $S_0(A, \mathbf{b})$ of *proper* solutions is the subset of $S(A, \mathbf{b})$ with all coordinates being pairwise distinct, that is

$$S_0(A, \mathbf{b}) = \{(x_1, \dots, x_m) \in S(A, \mathbf{b}) : x_i \neq x_j \text{ for all } 1 \leq i < j \leq m\}.$$

Before proceeding to the statement of our first main result, we need to define one more parameter, which will measure the densest subsystem of A . The motivation behind this can be compared to the graph setting, where one wants to study the number of occurrences of some fixed graph G as a subgraph of a binomial random graph with n vertices and edge probability p . Here, one naively would expect the threshold for a random graph to contain G to be when p is around $n^{-v(G)/e(G)}$, which is when the expected number of copies of G flips from 0 to positive. But this is not the case: if G contains a subgraph H with $v(H)/e(H) > v(G)/e(G)$, then $n^{-v(H)/e(H)}$ will define the threshold instead. A similar behavior occurs for the distribution of subgraph counts, as shown by Ruciński in [14].

Definition 3. For positive integers $m > r$ and a positive integer matrix $A \in \mathbb{Z}^{r \times m}$, define the *density* $c(A)$ of A by

$$c(A) = \max_{\emptyset \neq Q \subset [m]} \frac{|Q|}{|Q| - r_Q},$$

where $r_Q = r_Q(A) = \text{rk}(A) - \text{rk}(A^{\bar{Q}})$, and $\bar{Q} = [m] \setminus Q$.

Note that this is indeed well-defined, since for a positive matrix we see that $\text{rk}(A^{\bar{Q}}) \geq \text{rk}(A) - |Q| + 1$ for every $\emptyset \neq Q \subset m$ (see, for instance, the proof of Lemma 12).

We are now ready to state our first main theorem. For a random variable X with finite first moment $\mathbb{E}(X)$ and non-zero finite variance $\mathbb{V}\text{ar}(X)$, denote by $\tilde{X} = (X - \mathbb{E}(X))/\sqrt{\mathbb{V}\text{ar}(X)}$ its normalization. We write $X_n \xrightarrow{d} Y$ when a sequence of random variables $\{X_n\}_{n \geq 1}$ tends in distribution to Y .

Theorem 4. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ a positive and abundant integer matrix, and $\mathbf{b} \in \mathbb{Z}^r$ such that $S(A, \mathbf{b}) \neq \emptyset$. Furthermore, let n be an integer, $0 \leq p := p(n) \leq 1$ and X_n the random variable equal $|S_0(A, \mathbf{b}) \cap [n]_p^m|$, which counts the number of proper solutions $\mathbf{x} \in [n]_p^m$ to $A\mathbf{x} = \mathbf{b}$. Then

$$\tilde{X}_n \xrightarrow{d} \mathcal{N}(0, 1)$$

if $n(1 - p) \rightarrow \infty$ and $np^{c(A)} \rightarrow \infty$.

Note that in the specific case of k -APs, this problem was already investigated by Barhoumi-Andréani, Koch and Liu [15]. In fact, they prove not only results on the limiting distribution of the number of k -term arithmetic progressions in $[n]_p$ even when k is unbounded (but sub-logarithmic), they also establish a bivariate central limit theorem for the joint distribution when considering the counts of two distinct progression lengths. In contrast, Theorem 4 in the setting of progressions requires the length to be fixed.

Non-trivial solutions, and main theorem for their distribution Proper solutions are always of interest in the number theoretical context, but in a wide variety of situations we need to take care of non-proper solutions. For example, for $A = \begin{pmatrix} 1 & 1 & -1 & -1 \end{pmatrix}$, we see that $S(A, \mathbf{0})$ is the number of additive quadruples satisfying $a + b = c + d$. Clearly, if $\{a, b\} \cap \{c, d\} \neq \emptyset$, then these two sets must in fact be the same. On the other hand, solutions of the form $2a = c + d$ with $a \notin \{c, d\}$ are not proper but obviously of interest and might be considered as valid ones.

Keeping this in mind, we will now recall the notion of *non-trivial* solutions for systems of linear equations due to Rué, Spiegel and Zumalacárregui [13] which generalized an earlier notion of this for single line equations introduced by Ruzsa [16]. For a solution $\mathbf{x} = (x_1, \dots, x_m) \in S(A, \mathbf{b})$, we define $\mathfrak{p}(\mathbf{x}) \subset 2^{[m]}$ to be the partition of $[m]$ such that for any $i, j \in [m]$ it holds that $x_i = x_j$ if and only if i and j are in the same partition class of $\mathfrak{p}(\mathbf{x})$. One can view $\mathfrak{p}(\mathbf{x})$ as an ordered $|\mathfrak{p}(\mathbf{x})|$ -tuple $(C_1, \dots, C_{|\mathfrak{p}(\mathbf{x})|})$ such that $\min C_i < \min C_j$ whenever $i < j$. Doing this, we can now define the matrix $A_{\mathfrak{p}(\mathbf{x})}$ in the following way. Suppose the columns of A are denoted by $\mathbf{c}_1, \dots, \mathbf{c}_m \in \mathbb{Z}^r$, then

$$A_{\mathfrak{p}(\mathbf{x})} = \left(\sum_{i \in C_1} \mathbf{c}_i \mid \sum_{i \in C_2} \mathbf{c}_i \mid \cdots \mid \sum_{i \in C_{|\mathfrak{p}(\mathbf{x})|}} \mathbf{c}_i \right),$$

so $A_{\mathfrak{p}(\mathbf{x})} \in \mathbb{Z}^{r \times |\mathfrak{p}(\mathbf{x})|}$ is the $r \times |\mathfrak{p}(\mathbf{x})|$ matrix obtained by contracting all columns in the same partition class, with columns ordered by the minimum index in each class. Finally, define the set

$$\mathfrak{P}(A) = \{ \mathfrak{p} \subset 2^{[m]} : \mathfrak{p} \text{ is a partition of } [m] \text{ and } \text{rk}(A_{\mathfrak{p}}) = \text{rk}(A) \}.$$

We are finally ready to introduce the notion of a *non-trivial* solution.

Definition 5. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, and $\mathbf{b} \in \mathbb{Z}^r$. Then the set $S_1(A, \mathbf{b})$ of *non-trivial* solutions is the subset of $S(A, \mathbf{b})$ with associated partitions coming from $\mathfrak{P}(A)$, that is

$$S_1(A, \mathbf{b}) = \{ \mathbf{x} \in S(A, \mathbf{b}) : \mathfrak{p}(\mathbf{x}) \in \mathfrak{P}(A) \}.$$

This definition might seem quite arbitrary, but the interested reader is invited to read the discussions in [13] and [12] which show that, in some sense, it is quite natural in

that it encompasses the natural notions of non-triviality for specific systems of linear equations studied in the literature. When we want to investigate the distribution of nontrivial solutions, we actually need to look at $c(A_{\mathbf{p}})$ for any partition type \mathbf{p} that is to be considered. A special case in which it suffices to only consider $c(A)$ is that of *strictly balanced* systems.

Definition 6. Let $m > r$ be positive integers and $A \in \mathbb{Z}^{r \times m}$ be positive. Then the matrix A is called *strictly balanced* if

$$c(A) = \frac{m}{m - \text{rk}(A)} > \max_{\emptyset \subsetneq Q \subsetneq [m]} \frac{|Q|}{|Q| - r_Q(A)},$$

and furthermore for every $\mathbf{p} \in \mathfrak{P}(A) \setminus \{\{1\}, \{2\}, \dots, \{m\}\}$ such that $A_{\mathbf{p}}$ is positive, it holds that $c(A) > c(A_{\mathbf{p}})$.

We can prove the following theorem about the distribution of nontrivial solutions of strictly balanced systems of linear equations.

Theorem 7. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ a positive and abundant strictly balanced integer matrix such that $S(A, \mathbf{0}) \neq \emptyset$. Furthermore, let n be an integer, $0 \leq p := p(n) \leq 1$ and X_n the random variable $|S_1(A, \mathbf{0}) \cap [n]_p^m|$ that counts the number of non-trivial solutions $\mathbf{x} \in [n]_p^m$ to $A\mathbf{x} = \mathbf{0}$. Then

$$\tilde{X}_n \xrightarrow{d} \mathcal{N}(0, 1)$$

if $n(1 - p) \rightarrow \infty$ and $np^{c(A)} \rightarrow \infty$.

Both of our main results can be compared to those proved in [13] and [12]. In the former the authors investigated the threshold behavior of $|S_1(A, \mathbf{0}) \cap [n]_p^m|$ when $A \in \mathbb{Z}^{r \times m}$ is a positive matrix and in the latter, Spiegel extended this to $|S_1(A, \mathbf{b}) \cap [n]_p^m|$ for an arbitrary $\mathbf{b} \in \mathbb{Z}^r$ when A is positive and also abundant.

In that sense, their results hold in a more general setting, even when compared to our main technical result, Theorem 18. However, we note that the study of the distribution is more delicate than the threshold behavior, as can be seen by another result from [13], in which the authors explore the behavior at the threshold of nontrivial solutions to $A\mathbf{x} = \mathbf{0}$ for strictly balanced systems, that is, they work exactly in the same setting as in Theorem 7.

To conclude this section, let us mention that the core of the proof is based on a combination of algebraic ideas dealing with the matrix associated with the linear system of equations and the method of moments in probability theory to show convergence towards a normal distribution. The second part is highly inspired by the techniques developed by Ruciński in [14] where he proved the following result on the distribution

of the number of occurrences of a fixed graph G as a subgraph in a binomial random graph.

Theorem 8 ([14]). *Let G be a graph on $v(G)$ vertices and $e(G)$ edges. Furthermore, let n be an integer, $0 \leq p := p(n) \leq 1$ and X_n the random variable that the number of subgraphs of $G(n, p)$ that are isomorphic to G . Then*

$$\tilde{X}_n \xrightarrow{d} \mathcal{N}(0, 1)$$

if and only if $n^2(1 - p) \rightarrow \infty$ and $np^{d(G)} \rightarrow \infty$, where $d(G) = \max_{H \subset G} e(H)/v(H)$ is the density of G .

The proof of our main result, Theorem 18 which in particular implies Theorems 4 and 7 follows a similar structure as that of Theorem 8. Specifically, in order to show that the stated conditions on p are sufficient, one considers three different ranges of p and determines the exact structure of the objects that actually affect the moments.

It turns out that when viewed through the lens of hypergraphs, these important objects will actually be essentially the same for both graphs and systems of linear equations. However, we want to stress that actually establishing this fact requires the delicate study of the more difficult linear algebraic structure of the patterns we want to take into account, and hence our work highly differs from [14].

Plan of the paper In Section 2 we will state some needed prior results from [13] and [12], as well as prove intermediate lemmas that are needed to establish a meta result, namely Theorem 18 from which one can deduce Theorems 4 and 7. Section 3 will then be devoted entirely to stating and proving this main meta result, Theorem 18. To do this, we will be using the so-called method of moments to analyze three distinct regimes of the probability p . We conclude with a discussion on some remaining open problems in Section 4. In particular, we will investigate the question of whether the sufficient conditions on p in Theorems 4, 7 and 18 are also necessary.

2. Algebraic properties of systems of linear equations

Let us start by investigating the relation between proper and non-trivial solutions. For any $A \in \mathbb{Z}^{r \times m}$ and $\mathbf{b} \in \mathbb{Z}^r$ it is clear that

$$S_0(A, \mathbf{b}) \subset S_1(A, \mathbf{b}) \subset S(A, \mathbf{b}).$$

Furthermore, suppose $s > 1$ and we fix a specific $(C_1, \dots, C_s) = \mathbf{p} \in \mathfrak{P}(A)$ such that $\min C_i < \min C_j$ whenever $i < j$. Then if $\mathbf{x} = (x_1, \dots, x_m) \in S_1(A, \mathbf{b})$ satisfies $\mathbf{p}(\mathbf{x}) = \mathbf{p}$, we see that $\tilde{\mathbf{x}} = (x_{\min C_1}, \dots, x_{\min C_s}) \in \mathbb{Z}^s$ has s pairwise distinct coordinates and $A_{\mathbf{p}}\tilde{\mathbf{x}} = \mathbf{b}$, so $\tilde{\mathbf{x}} \in S_0(A_{\mathbf{p}}, \mathbf{b})$. Conversely, for any proper solution $\mathbf{x} = (x_1, \dots, x_s)$ in $S_0(A_{\mathbf{p}}, \mathbf{b})$, we see that the vector $\tilde{\mathbf{x}} \in \mathbb{Z}^m$ with i -th coordinate \tilde{x}_i defined as

$$\tilde{x}_i = \sum_{j=1}^s x_j \cdot \delta_{i \in C_j}$$

satisfies $\mathfrak{p}(\tilde{\mathbf{x}}) = \mathfrak{p}$ and $A\tilde{\mathbf{x}} = \mathbf{b}$ (here $\delta_{i \in C_j}$ denotes the indicator function which takes values 0 or 1 if $i \notin C_j$ and $i \in C_j$, respectively). So we get the following lemma, which is essentially Lemma 1.10 in [12].

Lemma 9 ([12]). *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, and $\mathbf{b} \in \mathbb{Z}^r$. Then for any partition $\mathfrak{p} \in \mathfrak{P}(A)$ and any subset $T \subset \mathbb{Z}$, it holds that*

$$|\{\mathbf{x} \in S_1(A, \mathbf{b}) : \mathfrak{p}(\mathbf{x}) = \mathfrak{p}\} \cap T^m| = |S_0(A_{\mathfrak{p}}, \mathbf{b}) \cap T^{|\mathfrak{p}|}|.$$

Lemma 9 will be very helpful because if A is positive, then the approximate size of $S_0(A, \mathbf{b})$ is not difficult to determine for any vector \mathbf{b} . Specifically, we have the following result which is Lemma 1.4 in [12].

Lemma 10 ([12]). *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, and $\mathbf{b} \in \mathbb{Z}^r$ such that $S(A, \mathbf{b}) \neq \emptyset$. Then if $n \in \mathbb{N}$ it holds that*

$$|S_0(A, \mathbf{b}) \cap [n]^m| \leq |S(A, \mathbf{b}) \cap [n]^m| \leq n^{m-\text{rk}(A)}.$$

Furthermore, if A is positive we also have

$$|S_0(A, \mathbf{b}) \cap [n]^m| = \Omega(n^{m-\text{rk}(A)}).$$

In [13] the authors managed to find a threshold probability function for when non-trivial solutions to a homogeneous system of linear equations appear in $[n]_p$, but for our current purposes we need to make one more specification regarding types of solutions. The issue essentially lies in the fact that while A is positive, the same might not hold for $A_{\mathfrak{p}}$ for some $\mathfrak{p} \in \mathfrak{P}(A)$. Note that this can only happen in the case $\mathbf{b} \neq \mathbf{0}$, since clearly $S_0(A_{\mathfrak{p}}, \mathbf{0}) \neq \emptyset$ implies that $A_{\mathfrak{p}}$ is positive. As a problematic example, consider for instance the matrix $A = (1 \ 1 \ 1 \ 1 \ -1)$ which is positive since for instance $\mathbf{x} = (1, 2, 3, 4, 10) \in \mathbb{N}^5$ is a proper solution to the system $A\mathbf{x} = \mathbf{0}$. But taking $\mathbf{b} = 6$, we see that $\mathbf{y} = (1, 2, 3, 3, 3)$ satisfies $A\mathbf{y} = \mathbf{b}$. We have $\mathfrak{p}(\mathbf{y}) = (\{1\}, \{2\}, \{3, 4, 5\})$ and hence $A_{\mathfrak{p}(\mathbf{y})} = (1 \ 1 \ 1)$ which is clearly not positive.

The problem arises as follows: if $A_{\mathfrak{p}}$ is not positive but we still have a significant number of solutions of this type – this can for instance be achieved by using the previous example as a gadget in a larger system – the number of solutions in the random binomial set will depend heavily on whether or not a specific bounded number of elements are included, and so getting any good results on their distribution is unlikely.

Having stated this motivation, we first introduce the concept of *positive partitions* as the subset $\mathfrak{P}_0(A) \subset \mathfrak{P}(A)$ such that

$$\mathfrak{P}_0(A) = \{\mathfrak{p} \in \mathfrak{P}(A) : A_{\mathfrak{p}} \text{ is positive}\}.$$

For any $\mathfrak{P} \subset \mathfrak{P}(A)$ we can now define the concept of *type- \mathfrak{P}* solutions.

Definition 11. Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, $\mathbf{b} \in \mathbb{Z}^r$, and $\mathfrak{P} \subset \mathfrak{P}(A)$. Then the set $S_{\mathfrak{P}}(A, \mathbf{b})$ of *type- \mathfrak{P}* solutions is the subset of $S(A, \mathbf{b})$ with associated partitions coming from \mathfrak{P} , that is

$$S_{\mathfrak{P}}(A, \mathbf{b}) = \{\mathbf{x} \in S(A, \mathbf{b}) : \mathfrak{p}(\mathbf{x}) \in \mathfrak{P}\}.$$

If $\mathfrak{P} = \{\mathfrak{p}\}$, we will write $S_{\mathfrak{p}}(A, \mathbf{b})$ instead of $S_{\{\mathfrak{p}\}}(A, \mathbf{b})$.

Non-trivial solutions are therefore the same as *type- $\mathfrak{P}(A)$* solutions. The following is a straightforward consequence of Lemma 9 and the upper bound in Lemma 10.

Lemma 12. Let $m > r$ and n be positive integers, $A \in \mathbb{Z}^{r \times m}$ positive, $\mathbf{b} \in \mathbb{Z}^r$, $\mathfrak{P} \subset \mathfrak{P}_0(A)$, and $Z \subset [n]$ a fixed set. Then

$$|\{\mathbf{y} \in S_{\mathfrak{P}}(A, \mathbf{b}) \cap [n]^m : \{\mathbf{y}\} \cap Z \neq \emptyset\}| = O(n^{m-\text{rk}(A)-1}).$$

Furthermore, if A is also abundant, then

$$|\{\mathbf{y} \in S_{\mathfrak{P}}(A, \mathbf{b}) \cap [n]^m : |\{\mathbf{y}\} \cap Z| \geq 2\}| = O(n^{m-\text{rk}(A)-2}).$$

Proof. We will consider each partition $\mathfrak{p} \in \mathfrak{P}$ separately. Due to Lemma 9 we can consider $\mathbf{y} \in S_0(A_{\mathfrak{p}}, \mathbf{b}) \cap [n]^{|\mathfrak{p}|}$. If \mathbf{y} contains an element of Z , then there exists some non-empty index set $Q \subset [|\mathfrak{p}|]$ and a vector $\mathbf{z} \in \mathbb{Z}^{|Q|}$ such that $\mathbf{y}^{\bar{Q}} \in S_0(A_{\mathfrak{p}}^{\bar{Q}}, \mathbf{b} - A_{\mathfrak{p}}^Q \mathbf{z}) \cap [n]^{|\mathfrak{p}|-|Q|}$ (note that here, $\bar{Q} = [|\mathfrak{p}|] \setminus Q$). By Lemma 10, we have

$$|S_0(A_{\mathfrak{p}}^{\bar{Q}}, \mathbf{b} - A_{\mathfrak{p}}^Q \mathbf{z}) \cap [n]^{|\mathfrak{p}|-|Q|}| \leq n^{|\mathfrak{p}|-|Q|-\text{rk}(A_{\mathfrak{p}}^{\bar{Q}})}.$$

$A_{\mathfrak{p}}$ is positive, which implies that

$$\text{rk}(A_{\mathfrak{p}}^{\bar{Q}}) \geq \text{rk}(A_{\mathfrak{p}}) - |Q| + 1 = \text{rk}(A) - |Q| + 1.$$

Indeed, the positivity implies that the removal of any single column will not change the rank of the resulting matrix, and since the rank of a matrix is the dimension of its column space, any subsequent removal decreases the rank by at most 1. Hence

$$|\mathfrak{p}| - |Q| - \text{rk}(A_{\mathfrak{p}}^{\bar{Q}}) \leq |\mathfrak{p}| - \text{rk}(A) - 1 \leq m - \text{rk}(A) - 1. \tag{1}$$

Since there are only $O(1)$ choices for \mathfrak{p} , Q and \mathbf{z} , this implies the result. For the second part, note first that for any $|\mathfrak{p}| < m$, Relation (1) actually implies the statement already,

so it only remains to consider the case $|\mathbf{p}| = m$, that is, $\mathbf{y} \in S_0(A, \mathbf{b}) \cap [n]^m$. If A is abundant, the removal of any two columns keeps the rank constant, while any subsequent removal will decrease it by at most 1 at a time. Hence for any $Q \subset [m]$ with $|Q| \geq 2$ we have

$$\text{rk}(A^{\bar{Q}}) \geq \text{rk}(A) - |Q| + 2,$$

and so

$$m - |Q| - \text{rk}(A^{\bar{Q}}) \leq m - \text{rk}(A) - 2,$$

which is what we wanted to show. \square

2.1. Compounded matrices

While Lemma 12 showed that the upper bound of Lemma 10 is always helpful in the situation of counting solutions with some entries fixed beforehand, the requirement of positivity is sometimes too restrictive to do the same for the lower bound. Suppose $\mathbf{x} \in S_1(A, \mathbf{b})$ is some fixed non-trivial solution. We want then to count the number of solutions $\mathbf{y} \in S_1(A, \mathbf{b})$ that intersect \mathbf{x} , that is, the size of the set $S_0(A^{\bar{Q}}_{\mathbf{p}(\mathbf{y})}, \mathbf{b} - A^{\bar{Q}}_{\mathbf{p}(\mathbf{y})}\tilde{\mathbf{x}})$, where $Q \subset [|\mathbf{p}(\mathbf{y})|]$ is some index set and $\tilde{\mathbf{x}} \in \{\mathbf{x}\}^{|Q|}$. Lemma 12 with $Z = \{\mathbf{x}\}$ immediately gives helpful upper bounds, but there are two issues when trying to apply the lower bound of Lemma 10 directly.

The first is that when summing over several distinct Q , the same solution will be counted multiple times, but this could be alleviated by just counting a single Q that maximizes the cardinality of the corresponding set. The bigger issue is that it is not clear at all that the matrix $A^{\bar{Q}}_{\mathbf{p}(\mathbf{y})}$ will be positive, and in fact this is not true in general even when $A_{\mathbf{p}(\mathbf{y})}$ is itself positive. Consider for instance the matrix $A = \begin{pmatrix} 1 & 1 & -1 \end{pmatrix}$ (associated with the Schur equation $x + y = z$), which is both positive and abundant. Then for any Q with $|Q| = 1$, the equation $A^{\bar{Q}}(x, y) = \mathbf{0}$ implies either $x = y$ or $x = -y$, so in either case positivity will be violated.

Instead, we will consider the concept of the *compounded* matrix already used in [13] to study the distribution at the threshold. For matrices $A \in \mathbb{Z}^{r_A \times m_A}$, $B \in \mathbb{Z}^{r_B \times m_B}$ and a bijection $M: P \rightarrow [m_B]$ with $P = \{p_1 < \dots < p_{|P|}\} \subset [m_A]$, define the $(r_A + r_B) \times (m_A + m_B - |P|)$ matrix $A \overset{M}{\times} B$ as

$$A \overset{M}{\times} B = \left(\begin{array}{c|c|c|c|c} A^{[m_A] \setminus P} & \mathbf{a}_{p_1} & \dots & \mathbf{a}_{p_{|P|}} & \mathbf{0} \\ \hline \mathbf{0} & \mathbf{b}_{M(p_1)} & \dots & \mathbf{b}_{M(p_{|P|})} & B^{[m_B] \setminus M(P)} \end{array} \right),$$

where $\mathbf{a}_1, \dots, \mathbf{a}_{m_A}$ denote the columns of A and $\mathbf{b}_1, \dots, \mathbf{b}_{m_B}$ the columns of B . Note that while we can apply this operator iteratively, the operation is in general not associative or even well-defined, that is

$$(A \times^M B) \times^{M'} C \neq A \times^M (B \times^{M'} C).$$

An exception to this is the case when $\text{dom}(M) = \emptyset$, which we will abbreviate by writing $A \times B$. In general, whenever no parentheses are used the compounded matrix is implied to be constructed from left to right iteratively, that is $A \times^M B \times^{M'} C := (A \times^M B) \times^{M'} C$. An important example that will appear often concerns the case $A = B$ and $M = \text{id}_Q$ the identity function for some index set $Q \subset [m_A]$, that is,

$$A \times^{\text{id}_Q} A = \left(\begin{array}{c|c|c} A^{\bar{Q}} & A^Q & \mathbf{0} \\ \hline \mathbf{0} & A^Q & A^{\bar{Q}} \end{array} \right).$$

The reason to study these matrices is natural: Let $\rho_A = \rho_A(M) : [m_A] \rightarrow [m_A]$ be the bijective function that maps the column indices of the first $[m_A]$ columns of $A \times^M A$ to the corresponding ones of A , and define $\rho_B : [m_A + m_B - |P|] \setminus [m_A - |P|] \rightarrow [m_B]$ similarly. Recall that for any integer $t > 1$, \mathbf{b}^t denotes the vector $(\mathbf{b}, \dots, \mathbf{b}) \in \mathbb{Z}^{rt}$. Then $\mathbf{z} = (z_1, \dots, z_{m_A+m_B-|P|})$ is a proper integer solution to the system of linear equations $(A \times^M B)\mathbf{z} = \mathbf{b}^2$ if and only if

$$\begin{aligned} \mathbf{x} &= (x_1, \dots, x_{m_A}) := (z_{\rho_A^{-1}(1)}, \dots, z_{\rho_A^{-1}(m_A)}) \text{ and} \\ \mathbf{y} &= (y_1, \dots, y_{m_B}) := (z_{\rho_B^{-1}(1)}, \dots, z_{\rho_B^{-1}(m_B)}) \end{aligned}$$

are proper solutions to the systems $A\mathbf{x} = \mathbf{b}$ and $B\mathbf{y} = \mathbf{b}$, and for any $i \in [m_A]$ and $j \in [m_B]$ it holds that $x_i = y_j$ if and only if $i \in P$ and $j = M(i)$. In other words, elements of $S_0(A \times^M B, \mathbf{b}^2)$ correspond to pairs of proper solutions in $S_0(A, \mathbf{b})$ and $S_0(B, \mathbf{b})$ that intersect exactly in the coordinates indicated by the function M .

We first state an easy but important property that any compounded matrix satisfies:

Lemma 13. *Let $m_A > r_A$ and $m_B > r_B$ be positive integers, $A \in \mathbb{Z}^{r_A \times m_A}$, $B \in \mathbb{Z}^{r_B \times m_B}$, $P \subset [m_A]$, $M : P \rightarrow [m_B]$ a bijection and $Q \subset [m_A + m_B - |P|]$. If we define $Q_1 = Q \cap [m_A]$ and $Q_2 = Q \setminus Q_1$, then*

$$\text{rk}((A \times^M B)^{\bar{Q}}) \geq \text{rk}(A^{[m_A] \setminus Q_1}) + \text{rk}(B^{\overline{M(P)} \setminus Q_2}).$$

In particular, $\text{rk}(A \times^M B) \geq \text{rk}(A) + \text{rk}(B^{\overline{M(P)}})$.

Proof. It is clear that the rank of $(A \times^M B)^{[m_A] \setminus Q_1}$ is at least $\text{rk}(A^{[m_A] \setminus Q_1}) =: r_1$, and similarly the rank of $(A \times^M B)^{[m_A+m_B-|P|] \setminus ([m_A] \cup Q_2)}$ is at least $\text{rk}(B^{\overline{M(P)} \setminus Q_2}) =: r_2$, so let $\mathbf{c}_1, \dots, \mathbf{c}_{r_1}$ be a collection of r_1 linearly independent vectors from the first $m_A - |Q_1|$ columns of $(A \times^M B)^{\bar{Q}}$ such that the corresponding column vectors of $A^{[m_A] \setminus Q_1}$ – that is,

the vectors obtained by only considering the first r_A rows – are also linearly independent, and $\mathbf{c}_{r_1+1}, \dots, \mathbf{c}_{r_1+r_2}$ a collection of r_2 of independent vectors from the $m_B - |P| - |Q_2|$ last columns. Then since the entries in the first r_A rows of $\mathbf{c}_{r_1+1}, \dots, \mathbf{c}_{r_1+r_2}$ are 0 and the column vectors obtained by only considering the first r_A rows of $\mathbf{c}_1, \dots, \mathbf{c}_{r_1}$ are linearly independent, any linear representation of $\mathbf{0}$

$$\lambda_1 \mathbf{c}_1 + \dots + \lambda_{r_1+r_2} \mathbf{c}_{r_1+r_2}$$

must satisfy $\lambda_1 = \dots = \lambda_{r_1} = 0$. But since the remaining r_2 columns were also linearly independent, we must have $\lambda_{r_1+1} = \dots = \lambda_{r_1+r_2} = 0$ as well, and hence the $r_1 + r_2$ columns $\mathbf{c}_1, \dots, \mathbf{c}_{r_1+r_2}$ are linearly independent. Since the rank of a matrix is the dimension of its column space we are done. \square

Next we will show that for certain compounded matrices this is indeed the correct rank.

Lemma 14. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$, and $Q \subset [m]$. Then*

$$\text{rk}(A \overset{\text{id}_Q}{\times} A) = \text{rk}(A) + \text{rk}(A^{\bar{Q}}).$$

Proof. We prove by obtaining the same upper and lower bounds for $\text{rk}(A \overset{\text{id}_Q}{\times} A)$. The lower bound follows from Lemma 13. For the upper bound, note that by performing elementary row and column operations we see that

$$\text{rk}(A \overset{\text{id}_Q}{\times} A) = \text{rk} \left(\begin{array}{c|c|c} A^Q & A^{\bar{Q}} & \mathbf{0} \\ \hline \mathbf{0} & -A^{\bar{Q}} & A^{\bar{Q}} \end{array} \right).$$

Let $r_1 = \text{rk}(A)$ and $r_2 = \text{rk}(A^{\bar{Q}})$ and suppose we have $r_1 + r_2 + 1$ row vectors from this matrix. If at least $r_1 + 1$ rows come from the upper half, then there exists a nontrivial representation of $\mathbf{0} \in \mathbb{Z}^{1 \times (2m - |Q|)}$, and hence the remaining coefficients can be set to 0. If on the other hand more than $r_2 + 1$ rows come from the bottom half, we can achieve a nontrivial representation of $\mathbf{0}$ only using these rows. Since one of those two cases must happen, we have the required upper bound. \square

The next result shows how positivity of A can (under some conditions) be passed on to a compounded matrix.

Lemma 15. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ be positive, and $Q \subset [m]$. If there exists an integer vector $\mathbf{y} \in S(A^{\bar{Q}}, \mathbf{0})$ with all coordinates nonzero, then $A \overset{\text{id}_Q}{\times} A$ is positive.*

Proof. Since A was positive, there exists a solution $\mathbf{x} = (x_1, \dots, x_m) \in S(A, \mathbf{0}) \cap \mathbb{N}^m$. But then if we define $\mathbf{y} = (\mathbf{x}^{\bar{Q}}, \mathbf{x}^Q, \mathbf{x}^{\bar{Q}})$, we clearly have $\mathbf{y} \in S(A \times^{\text{id}_Q} A, \mathbf{0}) \cap \mathbb{N}^m$. It remains to be proved that for any two distinct indices $i, j \in [2m - |Q|]$, there exists a solution $(x_1, \dots, x_{2m-|Q|}) = \mathbf{x}$ in $S(A \times^{\text{id}_Q} A, \mathbf{0})$ such that $x_i \neq x_j$. Let $\rho: [2m - |Q|] \rightarrow [m]$ denote the surjective function that maps the column indices of $A \times^{\text{id}_Q} A$ to the corresponding ones of A . If $\rho(i) \neq \rho(j)$, that is, i and j refer to distinct variables of A , this again just follows directly from the fact that A was positive itself. So we need to check the case $\rho(i) = \rho(j) \in \bar{Q}$. By our assumption on $A^{\bar{Q}}$, there exists a vector $\mathbf{y} \in S(A^{\bar{Q}}, \mathbf{0})$ with all its coordinates nonzero. Clearly, the same will hold for $2\mathbf{y}$, and hence

$$\mathbf{x} = (\mathbf{y}, \mathbf{0}, 2\mathbf{y}) \in \mathbb{Z}^{2m-|Q|}$$

will satisfy $(A \times^{\text{id}_Q} A)\mathbf{x} = \mathbf{0}$ and the i -th and j -th entry of \mathbf{x} are different. \square

Note that if $A \in \mathbb{Z}^{r \times m}$ is abundant, then any $Q \subset [m]$ of size $|Q| = 1$ will satisfy the requirements of Lemma 15. The next result shows that even in the case that Q itself does not satisfy them, we can find a superset that does, which will help us to get universal lower bounds that are sufficient for our applications.

Lemma 16. *Let $m > r$ be positive integers, $n \in \mathbb{N}$, $A \in \mathbb{Z}^{r \times m}$ positive, and $\mathbf{b} \in \mathbb{Z}^r$ such that $S(A, \mathbf{b}) \neq \emptyset$. Then for all $Q \subset [m]$ there exists $Q' \supset Q$ such that*

$$|S_0(A \times^{\text{id}_{Q'}} A, \mathbf{b}^2) \cap [n]^{2m-|Q'|}| = \Omega(n^{2m-\text{rk}(A)-\text{rk}(A^{\bar{Q}})-|Q'|}).$$

Proof. If Q satisfies the assumptions of Lemma 15, we see that $A \times^{\text{id}_Q} A$ is positive and hence we can apply the lower bound of Lemma 10 directly with $Q' = Q$, noting that $A \times^{\text{id}_Q} A$ has $2m - |Q|$ columns and rank $\text{rk}(A) + \text{rk}(A^{\bar{Q}})$ by Lemma 14. Otherwise, denote by $\mathbf{c}_1, \dots, \mathbf{c}_m \in \mathbb{Z}^r$ the columns of A , and let $Q_1 \subset \bar{Q}$ be the index set of the columns that will always have coefficient zero in a linear combination of the zero vector. We claim that these columns are linearly independent and their span does not contain any column \mathbf{c}_i with $i \in \bar{Q} \setminus Q_1$. Indeed, linear independence holds because any non-trivial linear combination $\sum_{i \in Q_1} \lambda_i \mathbf{c}_i = \mathbf{0}$ could be extended to a linear combination $\sum_{i \in \bar{Q}} \lambda_i \mathbf{c}_i = \mathbf{0}$ such that $\lambda_i \neq 0$ for at least one $i \in Q_1$, a contradiction to the definition of Q_1 . Similarly, if $\sum_{i \in Q_1} \lambda_i \mathbf{c}_i = \mathbf{c}_j$ for some $j \in \bar{Q} \setminus Q_1$, we obviously see that $\sum_{i \in Q_1} \lambda_i \mathbf{c}_i - \mathbf{c}_j = \mathbf{0}$ is a linear combination with at least one $\lambda_i \neq 0$, again a contradiction.

We thus see that defining $Q' = Q \cup Q_1$, the matrix $A^{\bar{Q}'}$ has rank $\text{rk}(A^{\bar{Q}'}) = \text{rk}(A^{\bar{Q}}) - |Q_1|$ and its number of columns is $m - |Q'| = m - |Q| - |Q_1|$, and hence

$$m - |Q'| - \text{rk}(A^{\bar{Q}'}) = m - |Q| - \text{rk}(A^{\bar{Q}}).$$

By construction, $A^{\overline{Q}'}$ satisfies the assumption of Lemma 15, and hence for $n \in \mathbb{N}$, applying Lemma 10 implies

$$|S_0(A \overset{\text{id}_{Q'}}{\times} A, \mathbf{b}^2) \cap [n]^{2m-|Q'|}| = \Omega(n^{2m-\text{rk}(A)-\text{rk}(A^{\overline{Q}})-|Q|}),$$

which is what we wanted to show. \square

Lemmas 14 and 15 actually apply to a more general iterated construction. Using the assumptions made in Lemma 13, write $Q = \{q_1 < \dots < q_{|Q|}\}$, and define for any integer $j \geq 1$ the bijection $M_j: [j(m - |Q|) + |Q|] \setminus [j(m - |Q|)] \rightarrow [|Q|]$ by $M_j(i) = q_{i-j(m-|Q|)}$. Then all the results mentioned in the aforementioned lemmas also apply in a natural way to the matrix

$$A \overset{\text{id}_Q}{\times} A \overset{M_1}{\times} \dots \overset{M_t}{\times} A = \begin{pmatrix} A^{\overline{Q}} & & & A^Q & & \\ & \ddots & & \vdots & & \\ & & A^{\overline{Q}} & A^Q & & \\ & & & A^Q & A^{\overline{Q}} & \end{pmatrix} \tag{2}$$

for any $t \geq 1$. The idea being that instead of just having a pair of proper solutions to $A\mathbf{x} = \mathbf{b}$, we now have a collection of $t + 1$ of them that all mutually intersect exactly in the variables indexed by Q . Specifically, we get the following result.

Lemma 17. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ be positive and abundant and $Q = \{i\} \subset [m]$ of size 1, and for any positive integer j let M_j denote the function $j m - j + 1 \mapsto i$. Then for any positive integer t it holds that the matrix $A \overset{\text{id}_Q}{\times} A \overset{M_1}{\times} \dots \overset{M_t}{\times} A$ is positive and of rank*

$$\text{rk}(A \overset{\text{id}_Q}{\times} A \overset{M_1}{\times} \dots \overset{M_t}{\times} A) = \text{rk}(A) + (t + 1) \text{rk}(A^{\overline{Q}}).$$

Proof. The proof is identical to that of Lemmas 14 and 15, noting that since A is abundant, for any $Q \subset [m]$ of size 1 it will hold that $S(A^{\overline{Q}}, \mathbf{0})$ contains a solution with all entries non-zero. \square

3. The distribution of type- \mathfrak{P} solutions: the main meta theorem

We now state our main result (which we call *meta theorem*) that will imply Theorems 4 and 7.

Theorem 18. *Let $m > r$ be positive integers, $A \in \mathbb{Z}^{r \times m}$ a positive and abundant integer matrix, $\mathbf{b} \in \mathbb{Z}^r$ such that $S(A, \mathbf{b}) \neq \emptyset$, and $\mathfrak{P} \subset \mathfrak{P}(A)$ satisfying $(\{1\}, \dots, \{m\}) \in \mathfrak{P}$ and*

$$\{\mathfrak{p} \in \mathfrak{P} : S_0(A_{\mathfrak{p}}, \mathbf{b}) \neq \emptyset\} \subset \mathfrak{P}_0(A).$$

Furthermore, let n be an integer, $0 \leq p := p(n) \leq 1$ and X_n the random variable $|S_{\mathfrak{P}}(A, \mathbf{b}) \cap [n]_p^m|$ that counts the number of type- \mathfrak{P} solutions $\mathbf{x} \in [n]_p^m$ to $A\mathbf{x} = \mathbf{b}$. Then

$$\tilde{X}_n \xrightarrow{d} \mathcal{N}(0, 1)$$

if $n(1 - p) \rightarrow \infty$ and $np^{c(A_p)} \rightarrow \infty$ for all $\mathfrak{p} \in \mathfrak{P}$ with $S_{\mathfrak{p}}(A, \mathbf{b}) \neq \emptyset$.

Let us first see that this indeed implies Theorems 4 and 7.

Proof of Theorem 4. We apply Theorem 18 with $\mathfrak{P} = \{(\{1\}, \dots, \{m\})\}$, noting that $S_0(A, \mathbf{b}) \neq \emptyset$ follows from the assumptions and Lemma 10. \square

Proof of Theorem 7. We see that for any $\mathfrak{p} \in \mathfrak{P}(A)$, whenever $S_0(A_{\mathfrak{p}}, \mathbf{0}) \neq \emptyset$, the matrix $A_{\mathfrak{p}}$ is positive by definition, that is $\mathfrak{p} \in \mathfrak{P}_0(A)$. Since A is strictly balanced, it holds that

$$c(A) = \frac{m}{m - \text{rk}(A)} > \max_{\emptyset \subsetneq Q \subsetneq [m]} \frac{|Q|}{|Q| - r_Q(A)} > c(A_{\mathfrak{p}})$$

for any $\mathfrak{p} \in \mathfrak{P}(A)$ with $|\mathfrak{p}| < m$ and hence $np^{c(A_{\mathfrak{p}})} \rightarrow \infty$. \square

The proof of Theorem 18 involves the analysis of several sub-cases, which we will split into different subsections. In general, it will follow the ideas that were used by Ruciński in [14] where he proved a similar result to Theorem 18 in order to determine the distribution of the number of occurrences of a fixed graph G as a subgraph in a binomial random graph.

3.1. The proof of Theorem 18

Before going into further case analysis, let us first discuss the common jumping off point and strategy. Let $\mu_k = \mathbb{E}((X_n - \mathbb{E}(X_n))^k)$ denote the k -th central moment of X_n associated to the system of linear equations $A\mathbf{x} = \mathbf{b}$. Our final goal will always be to show that, independently of the system studied,

$$\mu_k = \begin{cases} (1 + o(1)) \frac{k!}{(k/2)!} 2^{-k/2} \mu_2^{k/2} & \text{if } k \text{ is even,} \\ o(\mu_2^{k/2}) & \text{if } k \text{ is odd.} \end{cases} \tag{3}$$

These estimates would show that the moments $\mathbb{E}(\tilde{X}_n^k)$ converge to the moments of a normal distribution, which is uniquely determined by its moments.

Given a solution $\mathbf{x} = (x_1, \dots, x_m)$ in $S(A, \mathbf{b})$, denote by $\mathbb{I}_{\mathbf{x}}$ the indicator random variable for the event $\{x_1, \dots, x_m\} \subset [n]_p$. Abusing notation somewhat, if $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in S(A, \mathbf{b})^k$ is a k -tuple of solutions, we will write $\{\chi\} = \bigcup_{i=1}^k \{\mathbf{x}_i\}$. As a visual shorthand, bold Latin letters will indicate solutions, while greek letters will denote tuples of solutions. Note first that by definition we have

$$\mu_k = \sum_{\chi \in \mathfrak{S}_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \left(\mathbb{I}_{\mathbf{x}} - p^{|\{\mathbf{x}\}|} \right) \right), \tag{4}$$

where \mathfrak{S}_k is the set containing all k -tuples $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in (S_{\mathfrak{P}}(A, \mathbf{b}) \cap [n]^m)^k$ such that for every $i \in [k]$ there exists a $j \in [k] \setminus \{i\}$ such that $\{\mathbf{x}_i\} \cap \{\mathbf{x}_j\} \neq \emptyset$. Indeed, for any k -tuple χ that contains some $\mathbf{x} \in S_{\mathfrak{P}}(A, \mathbf{b}) \cap [n]^m$ that is disjoint from the remaining $k - 1$ solutions, we have

$$\mathbb{E} \left(\prod_{\mathbf{y} \in \chi} \left(\mathbb{I}_{\mathbf{y}} - p^{|\{\mathbf{y}\}|} \right) \right) = \mathbb{E} \left(\mathbb{I}_{\mathbf{x}} - p^{|\{\mathbf{x}\}|} \right) \mathbb{E} \left(\prod_{\substack{\mathbf{y} \in \chi \\ \mathbf{y} \neq \mathbf{x}}} \left(\mathbb{I}_{\mathbf{y}} - p^{|\{\mathbf{y}\}|} \right) \right) = 0.$$

Equation (4) behaves slightly different depending on the behavior of p , so suppose $p \rightarrow a$ for some constant $a \in [0, 1]$. We split the analysis in three cases, depending on whether this limit belongs to $(0, 1)$, is equal to 1 and is equal to 0.

3.1.1. Case 1: $0 < a < 1$

In this case we see that (4) implies $\mu_k = \sum_{\mathfrak{S}_k} \Theta(1)$, so we need to analyze the cardinality of \mathfrak{S}_k . We are going to prove (3) by induction on k . The base cases $k = 2$ and $k = 1$ are clearly always true. Our induction hypothesis tells us then that for any $\ell < k$ it holds that $\mu_\ell = O(\mu_2^{\ell/2})$ and hence by our previous observation

$$|\mathfrak{S}_\ell| = O(|\mathfrak{S}_2|^{\ell/2}).$$

Let us now study the statement for parameter k . The analysis will be different depending on whether k is odd or even. Let \mathfrak{S}'_k denote the subset of \mathfrak{S}_k containing all k -tuples of solutions $(\mathbf{x}_1, \dots, \mathbf{x}_k)$ with $\mathbf{x}_i \in (S_{\mathfrak{P}}(A, \mathbf{b}) \cap [n]^m)$ such that for every i the choice of $j \neq i$ with $\{\mathbf{x}_i\} \cap \{\mathbf{x}_j\} \neq \emptyset$ is unique, that is, the solutions can be grouped into $k/2$ pairwise disjoint pairs. Note that clearly, $\mathfrak{S}'_k = \emptyset$ for every odd k . If $\overline{\mathfrak{S}}'_k = \mathfrak{S}_k \setminus \mathfrak{S}'_k$, we will show that

$$|\overline{\mathfrak{S}}'_k| = o(|\mathfrak{S}_2|^{k/2}),$$

which would imply

$$\sum_{\chi \in \overline{\mathfrak{S}}'_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \left(\mathbb{I}_{\mathbf{x}} - p^{|\{\mathbf{x}\}|} \right) \right) = o(\mu_2^{k/2}).$$

To see that this is true, suppose $\chi := (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \overline{\mathfrak{S}}'_k$. Then there must exist an index $i = i(\chi) \in [k]$ such that

$$\chi^{[k] \setminus \{i\}} = (\mathbf{x}_1, \dots, \mathbf{x}_{i-1}, \mathbf{x}_{i+1}, \dots, \mathbf{x}_k)$$

is contained in \mathfrak{S}_{k-1} . Taking the minimum choice of i for each χ , we have thus defined a map $\pi: \overline{\mathfrak{S}}'_k \rightarrow \mathfrak{S}_{k-1}$, and hence

$$|\mathfrak{S}'_k| \leq \max_{v \in \mathfrak{S}_{k-1}} |\pi^{-1}(v)| \cdot |\mathfrak{S}_{k-1}|.$$

Since by induction $|\mathfrak{S}_{k-1}| = O(|\mathfrak{S}_2|^{(k-1)/2})$, it suffices to show that

$$\max_{v \in \mathfrak{S}_{k-1}} |\pi^{-1}(v)| = o(|\mathfrak{S}_2|^{1/2}).$$

We will first determine a lower bound on $|\mathfrak{S}_2|$. It is clear that $|\mathfrak{S}_2| \geq |S_0(A \overset{\text{id}_Q}{\times} A, \mathbf{b}^2) \cap [n]^{2m-|Q|}|$ for any $Q \subset [m]$. Since A is abundant, any Q containing exactly one index will satisfy the conditions of Lemma 15, so fix an arbitrary one. It follows by this and Lemma 10 applied to $A \overset{\text{id}_Q}{\times} A$ that

$$|\mathfrak{S}_2| \geq |S_0(A \overset{\text{id}_Q}{\times} A, \mathbf{b}^2) \cap [n]^{2m-1}| = \Omega(n^{2m-2 \text{rk}(A)-1}). \tag{5}$$

We now turn to giving an upper bound on $\max_{v \in \mathfrak{S}_{k-1}} |\pi^{-1}(v)|$, so fix an $v \in \mathfrak{S}_{k-1}$. By definition of \mathfrak{S}_k , any solution that could extend v to a k -tuple in \mathfrak{S}_k must intersect $\{v\}$, and hence by Lemma 12 we see that

$$|\pi^{-1}(v)| = O(n^{m-\text{rk}(A)-1}) = o(|\mathfrak{S}_2|^{1/2})$$

by the previously obtained lower bound (5). Since $\mathfrak{S}'_k = \emptyset$ for odd k , this also proves the odd case of (3).

We now turn to the case of even k . In this situation, the tuples that are summed over in μ_k are $k/2$ pairwise disjoint pairs of those summed over in μ_2 . We begin by showing that for all but a negligible amount, pairs $(\mathbf{x}, \mathbf{y}) \in \mathfrak{S}_2$ will satisfy $\mathbf{x}, \mathbf{y} \in S_0(A, \mathbf{b})$ and $|\{\mathbf{x}\} \cap \{\mathbf{y}\}| = 1$. Let us first see that we can restrict ourselves to pairs of proper solutions. This follows from a similar argument as was used in the proof of Lemma 12. Namely, if we fix an arbitrary $\mathbf{x} \in S_{\mathfrak{P}}(A, \mathbf{b})$, then by (1)

$$|\{y \in S_0(A_{\mathfrak{p}}, \mathbf{b}) \cap [n]^{|\mathfrak{p}}| : \mathfrak{p} \in \mathfrak{P}, |\mathfrak{p}| < m\}| = O(n^{m-\text{rk}(A)-2}).$$

Since there are at most $n^{m-\text{rk}(A)}$ choices for \mathbf{x} , this is negligible when compared to $|\mathfrak{S}_2|$. Applying the second part of Lemma 12 directly, the number of pairs of proper solutions that intersect in at least two elements is negligible as well. Note that for any pair (\mathbf{x}, \mathbf{y}) satisfying the structure described above, we have $\mathbb{E}((\mathbb{I}_{\mathbf{x}} - p^{\{\mathbf{x}\}})(\mathbb{I}_{\mathbf{y}} - p^{\{\mathbf{y}\}})) \sim (a^{2m-1}(1-a))$.

Finally, note that for any pair $\chi \in \mathfrak{S}_2$, the number of pairs $v \in \mathfrak{S}_2$ that share elements with χ is negligible: Clearly, any such v will consist of an $\mathbf{x} \in S_0(A, \mathbf{b})$ satisfying

$\{\mathbf{x}\} \cap \{\chi\} \neq \emptyset$ and a $\mathbf{y} \in S_0(A, \mathbf{b})$ satisfying $\{\mathbf{y}\} \cap \{\mathbf{x}\} \neq \emptyset$. Applying Lemma 12 again, the number of such pairs is $O(n^{2m-2\text{rk}(A)-2}) = o(|\mathfrak{S}_2|)$. We conclude

$$|\mathfrak{S}_k| \sim \binom{|\mathfrak{S}_2|/2}{k/2} k!$$

and hence, as p tends to $a \in (0, 1)$,

$$\mu_k \sim \binom{|\mathfrak{S}_2|/2}{k/2} k! (a^{2m-1}(1-a))^{k/2} \sim \frac{k!}{(k/2)!} 2^{-k/2} \mu_2^{k/2},$$

which is what we wanted to prove. \square

3.1.2. Case 2: $\mathbf{a} = \mathbf{1}$

Recall that p is chosen such that $n(1-p)$ tends to infinity. This property will be specially relevant in this case. We will instead look at the complements, so if \mathbf{x} is a solution to $A\mathbf{x} = \mathbf{b}$, define $q = 1 - p$ and $\bar{\mathbb{I}}_{\mathbf{x}} = 1 - \mathbb{I}_{\mathbf{x}}$. We see that $\nu_{\mathbf{x}} := \mathbb{E}(\bar{\mathbb{I}}_{\mathbf{x}}) = 1 - p^{|\{\mathbf{x}\}|} \sim |\{\mathbf{x}\}|q$ when n tends to infinity. Let $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \mathfrak{S}_k$ be a k -tuple of solutions in $S_{\mathfrak{S}}(A, \mathbf{b}) \cap [n]^m$. Using these definitions we see that (4) can be written as

$$\mu_k = \sum_{\chi \in \mathfrak{S}_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} (\mathbb{I}_{\mathbf{x}} - p^{|\{\mathbf{x}\}|}) \right) = (-1)^k \sum_{\chi \in \mathfrak{S}_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} (\bar{\mathbb{I}}_{\mathbf{x}} - \nu_{\mathbf{x}}) \right). \tag{6}$$

We associate then a set system $\mathcal{H} = \mathcal{H}(\chi)$ to the k -tuple χ as follows: the vertex set $V(\mathcal{H})$ is just $\{\chi\}$, while the edges are $E(\mathcal{H}) = \{\{\mathbf{x}\} : \mathbf{x} \in \chi\}$. Suppose the vertex cover number $\tau(\mathcal{H})$ of the hypergraph is s , that is, there exist s elements $a_1, \dots, a_s \in \{\chi\}$ such that removing them destroys all solutions in χ , and no collection of less than s elements in $[n]$ achieves this. Then a straightforward computation shows that

$$\mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \bar{\mathbb{I}}_{\mathbf{x}} \right) = \Theta(q^s).$$

Furthermore, it is clear that for any $t \in [k - 1]$, every $(k - t)$ -sub-collection of χ will require removal of at least $s - t$ elements to destroy all solutions in it, that is, for any $I \subset [k]$ with $|I| = k - t$ it holds that

$$\mathbb{E} \left(\prod_{i \in I} \bar{\mathbb{I}}_{\mathbf{x}_i} \prod_{j \notin I} \nu_{\mathbf{x}_j} \right) = \begin{cases} \Theta(q^s), & \text{if } \chi^I \in \mathfrak{S}_{k-t}, \\ 0, & \text{otherwise.} \end{cases}$$

Note here, we interpret χ as a k -vector in which every coordinate is itself an m -vector, so χ^I denotes the vector obtained by only keeping the solutions indexed by I . Putting this together, we see that (6) becomes

$$\mu_k = \sum_{s=1}^k \sum_{\chi \in \mathfrak{S}_{k,s}} \Theta(q^s), \tag{7}$$

where $\mathfrak{S}_{k,s}$ is the subset of k -tuples $\chi \in \mathfrak{S}_k$ such that $\tau(\mathcal{H}(\chi)) = s$. We will now show that for any fixed s , only χ of a certain structure contribute significantly to (7). For this, define by $\mathfrak{S}'_{k,s}$ the set of s -milky ways, which are $\chi \in \mathfrak{S}_{k,s}$ such that $\mathcal{H}(\chi)$ is the union of s disjoint components, with all edges in a component intersecting in a unique vertex. Note that this set is only non-empty for $s \leq \lfloor k/2 \rfloor$. Furthermore, note that each component corresponds to a matrix as described in (2), call it B_i , hence by considering the compounded matrix $B_1 \dot{\times} \dots \dot{\times} B_s$ we see that Lemma 17 together with Lemma 10 implies that for any k and $s \leq \lfloor k/2 \rfloor$

$$|\mathfrak{S}'_{k,s}| = \Omega(n^{k(m-\text{rk}(A)-1)+s}). \tag{8}$$

The combination of these two lemmas also gives a matching upper bound: For any fixed way of dividing up the k solutions into s components, there are at most $O(n^{k(m-\text{rk}(A)-1)+s})$ corresponding k -tuples. Since k and s are fixed, the number of these partitions is $O(1)$, and hence

$$|\mathfrak{S}'_{k,s}| = O(n^{k(m-\text{rk}(A)-1)+s}). \tag{9}$$

We will now show by induction on k that for any k and s ,

$$|\overline{\mathfrak{S}}'_{k,s}| := |\mathfrak{S}_{k,s} \setminus \mathfrak{S}'_{k,s}| = O(n^{k(m-\text{rk}(A)-1)+s-1}).$$

This holds trivially for $k = 1$, and for $k = 2$, it follows immediately from Lemma 12, since A is abundant and $s = 1$ being the only vertex cover number leading to a non-empty set. So suppose $k \geq 3$ and note that the bound (9) together with the induction hypothesis in particular implies $|\mathfrak{S}_{\ell,s}| = O(n^{\ell(m-\text{rk}(A)-1)+s})$ for any $\ell < k$.

We will split up $\overline{\mathfrak{S}}'_{k,s}$ even further, into disjoint sets $\overline{\mathfrak{S}}''_{k,s}$, $\overline{\mathfrak{S}}'''_{k,s}$, and $\overline{\mathfrak{S}}''''_{k,s}$, which are defined as follows:

- a) The set $\overline{\mathfrak{S}}''_{k,s}$ will contain all $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \overline{\mathfrak{S}}'_{k,s}$ such that there are solutions $\mathbf{x}_i, \mathbf{x}_j \in \chi$ satisfying $\{\chi^{[k] \setminus \{i,j\}}\} \cap \{\chi^{\{i,j\}}\} = \emptyset$, and $|\{\mathbf{x}_i\} \cap \{\mathbf{x}_j\}| \geq 2$.
- b) The set $\overline{\mathfrak{S}}'''_{k,s}$ contains all $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \overline{\mathfrak{S}}'_{k,s} \setminus \overline{\mathfrak{S}}''_{k,s}$, such that there exists a solution $\mathbf{x}_i \in \chi$ satisfying $|\{\mathbf{x}\} \cap \{\chi^{[k] \setminus \{i\}}\}| \geq 2$ and $\chi^{[k] \setminus \{i\}} \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$.
- c) Finally, $\overline{\mathfrak{S}}''''_{k,s} = \overline{\mathfrak{S}}'_{k,s} \setminus (\overline{\mathfrak{S}}''_{k,s} \cup \overline{\mathfrak{S}}'''_{k,s})$ are the remaining non-milky ways.

We proceed to analyze the cardinality of each set.

- a) We start with case $\overline{\mathfrak{S}}''_{k,s}$. For every $\chi \in \overline{\mathfrak{S}}''_{k,s}$, there are indices i and j such that $\chi^{[k] \setminus \{i,j\}} \in \mathfrak{S}_{k-2,s-1}$, and taking the lexicographic smallest pair we have defined a map $\pi: \overline{\mathfrak{S}}''_{k,s} \rightarrow \mathfrak{S}_{k-2,s-1}$, which implies

$$\begin{aligned} |\overline{\mathfrak{S}}''_{k,s}| &\leq |\mathfrak{S}_{k-2,s-1}| \max_{v \in \mathfrak{S}_{k-2,s-1}} |\pi^{-1}(v)| \\ &= O\left(n^{(k-2)(m-\text{rk}(A)-1)+s-1} \max_{v \in \mathfrak{S}_{k-1,s}} |\pi^{-1}(v)|\right). \end{aligned}$$

Lemmas 10 and 12 imply that for any $v \in \mathfrak{S}_{k-2,s-1}$ we have $|\pi^{-1}(v)| = O(n^{2m-2\text{rk}(A)-2})$, and hence

$$|\overline{\mathfrak{S}}''_{k,s}| = O(n^{k(m-\text{rk}(A)-1)+s-1}). \tag{10}$$

- b) We continue with the analysis of $\overline{\mathfrak{S}}'''_{k,s}$. In this case, taking the smallest possible index i , this again defines a map $\pi: \overline{\mathfrak{S}}'''_{k,s} \rightarrow \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$ and we see by induction hypothesis that

$$\begin{aligned} |\overline{\mathfrak{S}}'''_{k,s}| &\leq |\mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}| \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\ &= O\left(n^{(k-1)(m-\text{rk}(A)-1)+s} \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)|\right). \end{aligned}$$

Again, for any $v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$, we see that $|\pi^{-1}(v)|$ is at most the number of solutions $\mathbf{y} \in S_{\mathfrak{P}}(A, \mathbf{b}) \cap [n]^m$ that contain at least 2 elements from $\{v\}$, which is $O(n^{m-\text{rk}(A)-2})$ by Lemma 12, and hence

$$|\overline{\mathfrak{S}}'''_{k,s}| = O(n^{k(m-\text{rk}(A)-1)+s-1}). \tag{11}$$

- c) Finally, it remains to look at $\overline{\mathfrak{S}}''''_{k,s}$. If $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k)$ in $\overline{\mathfrak{S}}''''_{k,s}$, we claim that there must exist an index i such that $\chi^{[k] \setminus \{i\}} \in \mathfrak{S}_{k-1,s-1}$. In the sequel, we will only consider the components of $\mathcal{H}(\chi)$ that do not intersect in a unique vertex, of which there is at least one since $\chi \notin \mathfrak{S}'_{k,s}$. First note that it is clear that since $\chi \notin \mathfrak{S}'_{k,s} \cup \overline{\mathfrak{S}}''_{k,s}$, there must exist an index i such that $\chi^{[k] \setminus \{i\}} \in \mathfrak{S}_{k-1,s-1} \cup \mathfrak{S}_{k-1,s}$ or in other words, the remaining solutions are still intersecting. Since $\chi \notin \overline{\mathfrak{S}}'''_{k,s}$, for all of these indices, it must hold that $|\{\mathbf{x}_i\} \cap \{\chi^{[k] \setminus \{i\}}\}| = 1$. Finally, for at least one index i of the previously considered, it must actually hold that there is a unique j_i such that

$$\{\mathbf{x}_i\} \cap \{\chi^{[k] \setminus \{i\}}\} = \{\mathbf{x}_i\} \cap \{\mathbf{x}_{j_i}\}. \tag{12}$$

Indeed, since the components we consider are not sunflowers, there must exist an \mathbf{x}_ℓ that intersects the rest of its component in at least two points, and so the negation of (12) would imply $\chi \in \overline{\mathfrak{S}}''''_{k,s}$, since \mathbf{x}_ℓ would be a valid choice. We see that for any valid i satisfying (12) it holds that $\chi^{[k] \setminus \{i\}} \in \mathfrak{S}_{k-1,s-1}$. Having established this, we repeat the arguments already used, taking the minimal valid i and defining the appropriate function $\pi: \overline{\mathfrak{S}}''''_{k,s} \rightarrow \mathfrak{S}_{k-1,s-1}$ and conclude

$$\begin{aligned} |\overline{\mathfrak{S}}'''_{k,s}| &\leq |\mathfrak{S}_{k-1,s-1}| \max_{v \in \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\ &= O(n^{(k-1)(m-\text{rk}(A)-1)+s-1}) \max_{v \in \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\ &= O(n^{k(m-\text{rk}(A)-1)+s-1}). \end{aligned}$$

Together with (8), we see that (10), (11) and (12) imply that only s -milky ways contribute meaningfully when $s \leq \lfloor k/2 \rfloor$. Sadly, this bound is not quite strong enough when $s > \lfloor k/2 \rfloor$, since s -milky ways do not exist here. For this, we will prove by induction on k that for any k and $s \geq \lfloor k/2 \rfloor$ it holds that

$$|\mathfrak{S}_{k,s}| = O(n^{k(m-\text{rk}(A)-1)+\lfloor k/2 \rfloor}). \tag{13}$$

Again, the cases $k = 1$ and $k = 2$ follow from previous observations. Moreover, the statement is trivially true for any k and $s = \lfloor k/2 \rfloor$, so suppose $k \geq 3$ and $s > \lfloor k/2 \rfloor$. If $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \mathfrak{S}_{k,s}$, then there must exist a least index i such that $\chi^{[k] \setminus \{i\}} \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$, which allows us again to define a projection map $\pi: \mathfrak{S}_{k,s} \rightarrow \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}$. Since $s \geq \lfloor k/2 \rfloor + 1$ we have $s - 1 \geq \lfloor (k - 1)/2 \rfloor$, and hence by induction hypothesis we have

$$\begin{aligned} |\mathfrak{S}_{k,s}| &\leq |\mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}| \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\ &= O(n^{(k-1)(m-\text{rk}(A)-1)+\lfloor (k-1)/2 \rfloor}) \max_{v \in \mathfrak{S}_{k-1,s} \cup \mathfrak{S}_{k-1,s-1}} |\pi^{-1}(v)| \\ &= O(n^{k(m-\text{rk}(A)-1)+\lfloor (k-1)/2 \rfloor}), \end{aligned}$$

which in particular implies (13). Since $q \rightarrow 0$, we thus have

$$|\mathfrak{S}_{k,s}|q^s = o(|\mathfrak{S}_{k,\lfloor k/2 \rfloor}|q^{\lfloor k/2 \rfloor}),$$

and so the terms with $s > \lfloor k/2 \rfloor$ in (7) can be discarded. Now, if $\chi \in \mathfrak{S}'_{k,s}$ is a milky way, we see that

$$\mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \bar{\mathbb{I}}_{\mathbf{x}} \right) \sim q^s,$$

while for any $I \in [k]$ with $|I| \leq k - 1$ we have

$$\mathbb{E} \left(\prod_{i \in I} \bar{\mathbb{I}}_{\mathbf{x}_i} \prod_{j \notin I} \nu_{\mathbf{x}_j} \right) = O(q^{s+1}) = o(q^s),$$

since the removal of any single solution does not impact the cover number, while any subsequent removal can reduce it by at most one. Furthermore, by our previous observations we have already seen that for any k and $s \leq \lfloor k/2 \rfloor$

$$|\mathfrak{S}'_{k,s}|q^s = \Theta(n^{k(m-\text{rk}(A)-1)}(nq)^s),$$

so since $nq \rightarrow \infty$ by assumption only the $s = \lfloor k/2 \rfloor$ term contributes significantly. Hence (7) can be rewritten as

$$\mu_k \sim \sum_{s=1}^{\lfloor k/2 \rfloor} |\mathfrak{S}'_{k,s}|q^s \sim |\mathfrak{S}'_{k,\lfloor k/2 \rfloor}|q^{\lfloor k/2 \rfloor}. \tag{14}$$

This proves (3) for odd k , since

$$\mu_k = O(n^{k(m-\text{rk}(A)-1)}(nq)^{(k-1)/2}) = o(n^{k(m-\text{rk}(A)-1)}(nq)^{k/2}),$$

while

$$\mu_2^{k/2} = \Omega(n^{k(m-\text{rk}(A)-1)}(nq)^{k/2}).$$

For the even case, it is easy to see that one can repeat the argument from Case 1 to show that only those milky ways with all solutions being proper contribute meaningfully and that among those, the number of pairs in $\mathfrak{S}'_{2,1}$ that intersect another pair is negligible, so just like in that case we conclude

$$\mu_k \sim \binom{|\mathfrak{S}'_{2,1}|/2}{k/2} k!q^{k/2} \sim \frac{k!}{(k/2)!} 2^{-k/2} \mu_2^{k/2}. \quad \square$$

3.1.3. Case 3: $\mathbf{a} = \mathbf{0}$

Finally, in this case the crucial property of p that will be used is the fact that $np^{c(A_{\mathfrak{p}})}$ tends to infinity for every considered partition type $\mathfrak{p} \in \mathfrak{P}$ with non-empty solution set.

Let $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \mathfrak{S}_k$. Then for any $I \subsetneq [k]$ we see that

$$\mathbb{E} \left(\prod_{i \in I} \mathbb{I}_{\mathbf{x}_i} \prod_{j \in [k] \setminus I} p^{|\{\mathbf{x}_j\}|} \right) = p^{|\cup_{i \in I} \{\mathbf{x}_i\}| + \sum_{j \in [k] \setminus I} |\{\mathbf{x}_j\}|} = o \left(p^{|\cup_{i=1}^k \{\mathbf{x}_i\}|} \right),$$

since every \mathbf{x}_j with $j \notin I$ intersects at least one other solution. Since we clearly have $\mathbb{E}(\mathbb{I}_{\mathbf{x}_1} \dots \mathbb{I}_{\mathbf{x}_k}) = p^{|\cup_{i=1}^k \{\mathbf{x}_i\}|}$, it thus follows that

$$\mathbb{E} \left((\mathbb{I}_{\mathbf{x}_1} - p^{|\{\mathbf{x}_1\}|}) \dots (\mathbb{I}_{\mathbf{x}_k} - p^{|\{\mathbf{x}_k\}|}) \right) \sim p^{|\cup_{i=1}^k \{\mathbf{x}_i\}|}. \tag{15}$$

Furthermore, since $np^{c(A_{\mathfrak{p}})} \rightarrow \infty$ for any $\mathfrak{p} \in \mathfrak{P}$, we see that for any nonempty $Q \subset \llbracket \mathfrak{p} \rrbracket$, it holds that

$$n^{|Q| - r_Q(A_{\mathfrak{p}})} p^{|Q|} = \omega(1).$$

Let us prove (3) via induction on k , the cases $k = 1$ and $k = 2$ clearly being true. Note that in particular, the induction hypothesis implies that for any $\ell < k$ it holds that $\mu_\ell = O(\mu_2^{\ell/2})$. Let us decompose μ_k as

$$\mu_k \sim \sum_{\chi \in \mathfrak{S}_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \mathbb{I}_{\mathbf{x}} \right) = \sum_{\chi \in \mathfrak{S}'_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \mathbb{I}_{\mathbf{x}} \right) + B_k.$$

Recall that \mathfrak{S}'_k denoted the subset of \mathfrak{S}_k such that every $\mathbf{x} \in \chi$ had a unique partner $\mathbf{y} \in \chi$ and was disjoint from all other components. Our first step will be to show that

$$B_k = o(\mu_2^{k/2}),$$

which would in particular imply (3) in the case of odd k , since here $\mathfrak{S}'_k = \emptyset$. To see this, note that $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k) \in \mathfrak{S}_k \setminus \mathfrak{S}'_k$ implies that there must exist an index $i \in [k]$ such that $\chi^{[k] \setminus \{i\}} \in \mathfrak{S}_{k-1}$. Let i_χ denote the least index $i \in [k]$ for which this is true, then we can define the map $\pi : \mathfrak{S}_k \setminus \mathfrak{S}'_k \rightarrow \mathfrak{S}_{k-1}$ by $\pi(\chi) = \chi^{[k] \setminus \{i_\chi\}}$. Furthermore, let $Q_\chi \subset [|\mathbf{p}(\mathbf{x}_{i_\chi})|]$ denote the index set of all the components of \mathbf{x}_χ that are contained in $\{\pi(\chi)\}$. Then

$$\begin{aligned} B_k &= \sum_{\chi \in \mathfrak{S}_k \setminus \mathfrak{S}'_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \mathbb{I}_{\mathbf{x}} \right) \\ &= \sum_{\chi \in \mathfrak{S}_k \setminus \mathfrak{S}'_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \pi(\chi)} \mathbb{I}_{\mathbf{x}} \right) p^{|\mathbf{p}(\mathbf{x}_{i_\chi})| - |Q_\chi|} \\ &\leq \sum_{v \in \mathfrak{S}_{k-1}} \mathbb{E} \left(\prod_{\mathbf{y} \in v} \mathbb{I}_{\mathbf{y}} \right) \sum_{\mathfrak{p} \in \mathfrak{P}} \sum_{Q \subset [|\mathfrak{p}|]} \sum_{z \in \{v\}^{|Q|}} p^{|\mathfrak{p}| - |Q|} |S_0(A_{\mathfrak{p}}^{\bar{Q}}, b - A_{\mathfrak{p}}^Q z) \cap [n]^{|\mathfrak{p}| - |Q|}| \\ &\leq \mu_{k-1} \sum_{\mathfrak{p} \in \mathfrak{P}} \sum_{Q \subset [|\mathfrak{p}|]} n^{|\mathfrak{p}| - |Q| - \text{rk}(A_{\mathfrak{p}}^{\bar{Q}})} p^{|\mathfrak{p}| - |Q|} \\ &= \mu_{k-1} \sum_{\mathfrak{p} \in \mathfrak{P}} \sum_{Q \subset [|\mathfrak{p}|]} n^{|\mathfrak{p}| - \text{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))} p^{|\mathfrak{p}| - |Q|} \end{aligned} \tag{16}$$

Since $\mu_{k-1} = O(\mu_2^{(k-1)/2})$ it thus suffices to show that the remaining expression in (16) is $o(\mu_2^{1/2})$. But by Lemma 16 there exists a $Q' \supset Q$ such that, using $p \rightarrow 0$ we see that

$$\begin{aligned} \mu_2 &\geq p^{2|\mathfrak{p}| - |Q'|} |S_0(A_{\mathfrak{p}} \overset{\text{id}_{Q'}}{\times} A_{\mathfrak{p}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{2|\mathfrak{p}| - |Q'|}| \\ &\geq p^{2|\mathfrak{p}| - |Q|} |S_0(A_{\mathfrak{p}} \overset{\text{id}_{Q'}}{\times} A_{\mathfrak{p}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{2|\mathfrak{p}| - |Q'|}| \\ &= \Omega(n^{2|\mathfrak{p}| - 2 \text{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))} p^{2|\mathfrak{p}| - |Q|}), \end{aligned}$$

and hence

$$\mu_2^{1/2} = \Omega(n^{|\mathfrak{p}| - \text{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))/2} p^{|\mathfrak{p}| - |Q|/2}) = \omega(n^{|\mathfrak{p}| - \text{rk}(A) - (|Q| - r_Q(A_{\mathfrak{p}}))} p^{|\mathfrak{p}| - |Q|}).$$

As stated before, this finishes the proof in the case of odd k , so suppose k is even. If $\mathfrak{p}, \mathfrak{q} \in \mathfrak{P}$ and $M: P \rightarrow Q$ is a bijection between nonempty $P \subset [|\mathfrak{p}|]$ and $Q \subset [|\mathfrak{q}|]$, we will call the triple $(\mathfrak{p}, \mathfrak{q}, M)$ *leading* if

$$p^{|\mathfrak{p}| + |\mathfrak{q}| - |P|} |S_0(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{|\mathfrak{p}| + |\mathfrak{q}| - |P|}| = \Omega(\mu_2),$$

and hence

$$\mu_2 \sim \sum_{\substack{(\mathfrak{p}, \mathfrak{q}, M) \\ \text{leading}}} p^{|\mathfrak{p}| + |\mathfrak{q}| - |\text{dom}(M)|} |S_0(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{|\mathfrak{p}| + |\mathfrak{q}| - |\text{dom}(M)|}|.$$

Let us first make an observation that will be helpful later. Suppose $|\mathfrak{p}| \geq |\mathfrak{q}|$, and let $M: P \rightarrow Q$ be a bijection between some index sets $P \subset [|\mathfrak{p}|]$ and $Q \subset [|\mathfrak{q}|]$. Then by Lemma 16 there exists a $P' \supset P$ such that using $np \rightarrow \infty$ and $p \rightarrow 0$ we see

$$\begin{aligned} \mu_2 &= \Omega \left(p^{2|\mathfrak{p}| - |P'|} |S_0(A_{\mathfrak{p}} \overset{\text{id}_{P'}}{\times} A_{\mathfrak{p}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{2|\mathfrak{p}| - |P'|}| \right) \\ &= \Omega \left(p^{2|\mathfrak{p}| - |P|} |S_0(A_{\mathfrak{p}} \overset{\text{id}_{P'}}{\times} A_{\mathfrak{p}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{2|\mathfrak{p}| - |P|}| \right) \\ &= \Omega \left(n^{2|\mathfrak{p}| - 2 \text{rk}(A) - (|P| - r_P(A_{\mathfrak{p}}))} p^{2|\mathfrak{p}| - |P|} \right) \tag{17} \\ &= \Omega \left(n^{|\mathfrak{p}| + |\mathfrak{q}| - 2 \text{rk}(A) - (|P| - r_P(A_{\mathfrak{p}}))} p^{|\mathfrak{p}| + |\mathfrak{q}| - |P|} \right) \\ &= \Omega \left(p^{|\mathfrak{p}| + |\mathfrak{q}| - |P|} |S_0(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{2|\mathfrak{p}| - |P|}| \right). \end{aligned}$$

Here the last line followed from the fact that $\text{rk}(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}) \geq \text{rk}(A_{\mathfrak{q}}) + \text{rk}(A_{\mathfrak{p}}^{\overline{P}}) = 2 \text{rk}(A) - r_P(A_{\mathfrak{p}})$. Hence, $(\mathfrak{p}, \mathfrak{q}, M)$ can be a leading overlap only if

- i) $|\mathfrak{p}| = |\mathfrak{q}|$,
- ii) $r_P(A_{\mathfrak{p}}) = r_Q(A_{\mathfrak{q}})$,
- iii) $\text{rk}(A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}}) = 2 \text{rk}(A) - r_P(A_{\mathfrak{p}})$,
- iv) For any matrix $C \in \left\{ A_{\mathfrak{p}} \overset{\text{id}_P}{\times} A_{\mathfrak{p}}, A_{\mathfrak{q}} \overset{\text{id}_Q}{\times} A_{\mathfrak{q}}, A_{\mathfrak{p}} \overset{M}{\times} A_{\mathfrak{q}} \right\}$ it holds that

$$|S_0(C \cap [n]^{2|\mathfrak{p}| - |P|})| = \Theta(n^{2|\mathfrak{p}| - (|P| - r_P(A_{\mathfrak{p}}))}),$$

and

- v) $(\mathfrak{p}, \mathfrak{p}, \text{id}_P)$ and $(\mathfrak{q}, \mathfrak{q}, \text{id}_Q)$ are leading overlaps,

since otherwise some of the Ω terms in Equation (17) would turn into ‘ ω ’ terms. Note that these things in particular imply that there cannot exist any $P' \supsetneq P$ satisfying

$$|S_0(A_{\mathbf{p}} \times^{id_{P'}} A_{\mathbf{p}}, \mathbf{b}^2) \cap [n]^{2|\mathbf{p}|-|P'|}| = \Omega(n^{2|\mathbf{p}|-(|P|-r_P(A_{\mathbf{p}}))}),$$

since $p \rightarrow 0$ would then imply that $(\mathbf{p}, \mathbf{p}, id_P)$ is not a leading triple. By the proof of Lemma 16 this means that there does not exist any index $i \in [|\mathbf{p}|]$ such that $x_i = 0$ for every solution $\mathbf{x} \in S(A_{\mathbf{p}}, \mathbf{0})$. Before continuing, let us quickly try to understand the preceding statements. Essentially, one would naively hope that all leading triples are of the form $(\mathbf{p}, \mathbf{p}, id_P)$ for some $P \subset [|\mathbf{p}|]$, but this is too optimistic because it ignores the inherent symmetry of some systems of linear equations. For instance, it is clear that for $A = (1 \ 1 \ 1 \ -1)$, the partitions $\mathbf{p} = (\{1, 2\}, \{3\}, \{4\})$ and $\mathbf{q} = (\{1, 3\}, \{2\}, \{4\})$ are technically different, but essentially the same in the sense that there exists a bijection between $S_0(A_{\mathbf{p}}, \mathbf{b})$ and $S_0(A_{\mathbf{q}}, \mathbf{b})$.

We can now continue with the actual proof by defining $\mathfrak{S}''_k \subset \mathfrak{S}'_k$ to be the set of k -tuples $\chi = (\mathbf{x}_1, \dots, \mathbf{x}_k)$ such that whenever $1 \leq i < j \leq k$ and $|\{\mathbf{x}_i\} \cap \{\mathbf{x}_j\}| = s > 0$, it holds that $(\mathbf{p}(\mathbf{x}_i), \mathbf{p}(\mathbf{x}_j), M_{i,j})$ is a leading triple. Here $M_{i,j}$ is the bijection defining the incidences between the distinct components of \mathbf{x}_i and \mathbf{x}_j . We will show that

$$\sum_{\chi \in \mathfrak{S}'_k \setminus \mathfrak{S}''_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \mathbb{I}_{\mathbf{x}} \right) = o(\mu_2^{k/2}).$$

To see this, we will again define a map $\pi: \mathfrak{S}'_k \setminus \mathfrak{S}''_k \rightarrow \mathfrak{S}'_{k-2}$ in the following way. Among all $1 \leq i < j \leq k$ such that $(\mathbf{p}(\mathbf{x}_i), \mathbf{p}(\mathbf{x}_j), M_{i,j})$ is not a leading triple, let $\{i_\chi, j_\chi\}$ be the set minimizing $\min\{i, j\}$. Then we can define $\pi(\chi) = \chi^{[k] \setminus \{i_\chi, j_\chi\}}$ and see that

$$\begin{aligned} & \sum_{\chi \in \mathfrak{S}'_k \setminus \mathfrak{S}''_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \chi} \mathbb{I}_{\mathbf{x}} \right) \\ &= \sum_{\chi \in \mathfrak{S}'_k \setminus \mathfrak{S}''_k} \mathbb{E} \left(\prod_{\mathbf{x} \in \pi(\chi)} \mathbb{I}_{\mathbf{x}} \right) p^{|\mathbf{p}(\mathbf{x}_{i_\chi})| + |\mathbf{p}(\mathbf{x}_{j_\chi})| - |\text{dom}(M_{i_\chi, j_\chi})|} \\ &\leq \mu_{k-2} \sum_{\substack{(\mathbf{p}, \mathbf{q}, M) \\ \text{not leading}}} p^{|\mathbf{p}| + |\mathbf{q}| - |\text{dom}(M)|} |S_0(A_{\mathbf{p}} \times^M A_{\mathbf{q}}, (\mathbf{b}, \mathbf{b})) \cap [n]^{|\mathbf{p}| + |\mathbf{q}| - |\text{dom}(M)|}| \\ &= o(\mu_2^{k/2}), \end{aligned}$$

the last line following from the induction hypothesis $\mu_{k-2} = O(\mu_2^{(k-2)/2})$ and the definition of leading triples, noting that there are only $O(1)$ choices for \mathbf{p}, \mathbf{q} and M . Let us enumerate the set of leading triples by $(\mathbf{p}_1, \mathbf{q}_1, M_1), \dots, (\mathbf{p}_u, \mathbf{q}_u, M_u)$, and for any choice of $0 \leq i_1, \dots, i_u \leq k/2$ such that $\sum i_j = k/2$, define the matrix $A(i_1, \dots, i_u)$ by

$$A(i_1, \dots, i_u) = \underbrace{(A_{\mathbf{p}_1} \overset{M_1}{\times} A_{\mathbf{q}_1}) \dot{\times} \dots \dot{\times} (A_{\mathbf{p}_1} \overset{M_1}{\times} A_{\mathbf{q}_1}) \dot{\times} \dots}_{i_1 \text{ times}} \underbrace{\dot{\times} (A_{\mathbf{p}_u} \overset{M_u}{\times} A_{\mathbf{q}_u}) \dot{\times} \dots \dot{\times} (A_{\mathbf{p}_u} \overset{M_u}{\times} A_{\mathbf{q}_u})}_{i_u \text{ times}},$$

and write ℓ_j for the expression $|\mathbf{p}_j| + |\mathbf{q}_j| - |\text{dom}(M_j)|$. We also need the following result.

Claim. *For any integers $0 \leq i_1, \dots, i_u \leq k/2$ satisfying $\sum i_j = k/2$ it holds that*

$$|S_0(A(i_1, \dots, i_u), \mathbf{b}^k) \cap [n]^{i_1 \ell_1 + \dots + i_u \ell_u}| \sim \prod_{j \in [u]} |S_0(A_{\mathbf{p}_j} \overset{M_j}{\times} A_{\mathbf{q}_j}, \mathbf{b}^2) \cap [n]^{\ell_j}|^{i_j}.$$

Proof of the claim: Let $(\mathbf{p}, \mathbf{q}, M)$ be a leading triple with $P := \text{dom}(M)$. We will first show that $A_{\mathbf{p}}^{\overline{P}}$ does not contain any column c_i such that

$$\text{rk}(A_{\mathbf{p}}^{\overline{P} \setminus \{i\}}) = \text{rk}(A_{\mathbf{p}}^{\overline{P}}) - 1.$$

But this follows since for any such hypothetical column, every solution $\mathbf{x} \in S(A_{\mathbf{p}}^{\overline{P}}, \mathbf{0})$ would satisfy $x_i = 0$, and we have seen before that this cannot happen for any leading triple. We are now able to prove the claim. Consider any $k/2$ -tuple χ of solutions in $\times_{j=1}^u (S_0(A_{\mathbf{p}_j} \overset{M_j}{\times} A_{\mathbf{q}_j}, \mathbf{b}^2) \cap [n]^{\ell_j})^{i_j}$, of which there are $\prod_{j \in [u]} |S_0(A_{\mathbf{p}_j} \overset{M_j}{\times} A_{\mathbf{q}_j}, \mathbf{b}^2) \cap [n]^{\ell_j}|^{i_j}$ many. When these solutions are all pairwise disjoint, they in fact define a solution in $S_0(A(i_1, \dots, i_u), \mathbf{b}^k) \cap [n]^{\ell_1 i_1 + \dots + \ell_u i_u}$. Otherwise, there exist indices $s, t \in [u]$ and a bijection $M: P \rightarrow Q$ with $\emptyset \neq P \subset [\ell_s]$ and $\emptyset \neq Q \subset [\ell_t]$ such that some pair $\mathbf{x}, \mathbf{y} \in \chi$ defines a solution in $S_0((A_{\mathbf{p}_s} \overset{M_s}{\times} A_{\mathbf{q}_s}) \overset{M}{\times} (A_{\mathbf{p}_t} \overset{M_t}{\times} A_{\mathbf{q}_t}), \mathbf{b}^4) \cap [n]^{\ell_s + \ell_t - |P|}$. But if we define $Q_1 = Q \cap [|\mathbf{p}_t|]$ and $Q_2 = Q \cap [\ell_t] \setminus [|\mathbf{p}_t|]$, we see that

$$\begin{aligned} & \text{rk}((A_{\mathbf{p}_s} \overset{M_s}{\times} A_{\mathbf{q}_s}) \overset{M}{\times} (A_{\mathbf{p}_t} \overset{M_t}{\times} A_{\mathbf{q}_t})) \\ & \geq \text{rk}(A_{\mathbf{p}_s} \overset{M_s}{\times} A_{\mathbf{q}_s}) + \text{rk}((A_{\mathbf{p}_t} \overset{M_t}{\times} A_{\mathbf{q}_t})^{\overline{Q}}) \\ & \geq \text{rk}(A_{\mathbf{p}_s} \overset{M_s}{\times} A_{\mathbf{q}_s}) + \text{rk}(A_{\mathbf{p}_t}^{\overline{Q_1}}) + \text{rk}(A_{\mathbf{q}_t}^{\overline{\text{im}(M_t)} \setminus Q_2}) \\ & \geq \text{rk}(A_{\mathbf{p}_s} \overset{M_s}{\times} A_{\mathbf{q}_s}) + \text{rk}(A_{\mathbf{p}_t}) - |Q_1| + 1 + \text{rk}(A_{\mathbf{q}_t}^{\overline{\text{im}(M_t)}}) - |Q_2| + 1 \\ & = \text{rk}(A_{\mathbf{p}_s} \overset{M_s}{\times} A_{\mathbf{q}_s}) + 2 \text{rk}(A) - |Q| + 2 - r_{\text{im}(M_t)}(A_{\mathbf{q}_t}). \end{aligned}$$

Here, for the third inequality we have used the previously established fact that removing any one column from $A_{\mathbf{q}_t}^{\overline{\text{im}(M_t)}}$ will not reduce the rank. The same clearly holds for $A_{\mathbf{p}_t}$ since this matrix is positive.

We are almost done, since now we can use the upper bound from Lemma 10 and see that

$$\begin{aligned}
 & |S_0((A_{p_s} \overset{M_s}{\times} A_{q_s}) \overset{M}{\times} (A_{p_t} \overset{M_t}{\times} A_{q_t}), \mathbf{b}^4) \cap [n]^{\ell_s + \ell_t - |P|}| \\
 & \leq n^{\ell_s + \ell_t - |Q| - \text{rk}((A_{p_s} \overset{M_s}{\times} A_{q_s}) \overset{M}{\times} (A_{p_t} \overset{M_t}{\times} A_{q_t}))} \\
 & \leq n^{\ell_s + \ell_t - \text{rk}(A_{p_s} \overset{M_s}{\times} A_{q_s}) - 2 \text{rk}(A) + r_{\text{im}(M_t)}(A_{q_t})} \\
 & = o(|S_0(A_{p_s} \overset{M_s}{\times} A_{q_s}, \mathbf{b}^2) \cap [n]^{\ell_s}| \cdot |S_0(A_{p_t} \overset{M_t}{\times} A_{q_t}, \mathbf{b}^2) \cap [n]^{\ell_t}|),
 \end{aligned}$$

the last line following from the observations made from (17). Since there are only $O(1)$ many choices for s, t , and M , this proves the claim. \square

Using the claim, we can now conclude

$$\begin{aligned}
 \mu_k & \sim \frac{k!}{2^{k/2}} \sum_{\substack{i_1, \dots, i_u \\ i_1 + \dots + i_u = k/2}} \frac{1}{i_1! \dots i_u!} p^{i_1 \ell_1 + \dots + i_u \ell_u} |S_0(A(i_1, \dots, i_u), \mathbf{b}^k) \cap [n]^{i_1 \ell_1 + \dots + i_u \ell_u}| \\
 & \sim \frac{k!}{(k/2)! 2^{k/2}} \sum_{\substack{i_1, \dots, i_u \\ i_1 + \dots + i_u = k/2}} \binom{k/2}{i_1, \dots, i_u} \prod_{j \in [u]} (p^{\ell_j} |S_0(A_{p_j} \overset{M_j}{\times} A_{q_j}, \mathbf{b}^2) \cap [n]^{\ell_j}|)^{i_j} \\
 & \sim \frac{k!}{(k/2)! 2^{k/2}} \mu_2^{k/2},
 \end{aligned}$$

which is what we wanted to prove. \square

4. Concluding remarks

In this paper we have established sufficient conditions for the choice of p in order to guarantee normal limiting distributions for the number of solutions to linear systems of equations of the form $A\mathbf{x} = \mathbf{b}$. Specifically, we showed that $n(1 - p) \rightarrow \infty$ and $np^{c(A_p)} \rightarrow \infty$ (for all partitions under consideration) sufficed, and we used them in their full strength in Cases 2 and 3, respectively of the proof of Theorem 18. Both of these conditions are analogous to those that appear in Ruciński’s proof of normality for the number of copies of a given subgraph H in the binomial random model $G(n, p)$. Note that when comparing to the graph setting, the elements of $[n]$ in systems take on the function of both vertices and edges, and hence the analogue of our $n(1 - p) \rightarrow \infty$ requirement is to ask that $n^2(1 - p)$, the expected number of edges, is unbounded.

In fact, in [14] Ruciński showed that those conditions were necessary as well as sufficient. In our context we can say something similar regarding the condition that $n(1 - p)$ is unbounded. Observe that the expression $n(1 - p)$ can be interpreted as the expected number of elements that are *not* chosen in $[n]_p$, hence if $n(1 - p) \not\rightarrow \infty$, then $[n]_p$ is typically all the interval $[n]$ with the exception of a bounded number of elements. Then it is easy to show that under this condition, X_n is strongly concentrated around its mean value, concluding that $\tilde{X}_n \xrightarrow{d} 0$, and hence, we do not have a normal limiting distribution for the number of solutions.

However, the argument used by Ruciński in order to study the second condition requires a delicate study of moments of order 4 and 6, which we are unable to adapt. Roughly speaking, in our algebraic setting the structure of solutions is more complicated compared with the graph setting, as solutions with repeated components can be valid ones (the analogy would be to consider also subgraphs of the fixed graph H that we want to count the corresponding number of subcopies). Hence, an open question arising from our work is to obtain an *only if* statement of Theorem 18. As an intermediate step, one could investigate the two main corollaries of our meta theorem, Theorems 4 and 7, and in fact, the following arguments suggest that the sufficient conditions on p stated in these theorems are also necessary.

Let us start with the setting of Theorem 7, that of non-trivial solutions in strictly balanced homogeneous systems of linear equations. Here, as discussed before, Rué, Spiegel and Zumalacárregui in [13] already established a threshold result showing that when $np^{c(A)} = o(1)$, asymptotically almost surely $S_1(A, \mathbf{0}) \cap [n]_p^m$ is empty. Using a concentration argument similar to the one above when discussing the necessity of $n(1 - p) \rightarrow \infty$, we see that in this case $|S_1(A, \mathbf{0}) \cap [n]_p^m|$ will converge in distribution to the constant 0 distribution. In addition to this threshold results, the authors also studied the distribution at the threshold, that is, the case $np^{c(A)} \rightarrow a > 0$ where a is a constant. Their results show that here, the random variable $|S_1(A, \mathbf{0}) \cap [n]_p^m|$ converges in distribution to a Poisson. Putting all of this together, we see that this establishes the only if direction for Theorem 7.

Let us now turn to the setting of Theorem 4, that of proper solutions, where there are no repeated variables by definition. The argument that $np^{c(A)} = o(1)$ implies that $\tilde{X}_n \xrightarrow{d} 0$ is the same as before, so what remains to do is to understand the case when $np^{c(A)}$ tends to some positive constant. When A is strictly balanced, one can use the same arguments that were used by Rué, Spiegel and Zumalacárregui in the proof of the strictly balanced homogeneous case for the distribution of non-trivial solutions to see that one will also have a Poisson distribution in the setting of Theorem 4. When A is not strictly balanced, an analysis as was performed by Ruciński in the subgraph setting is needed, but since we are now in the situation that all variables are pairwise distinct, the same result should follow, which would establish the necessity of $np^{c(A)} \rightarrow \infty$.

To conclude, let us mention that once one has proved a normal limiting distribution, a natural next question is the study local limit theorems as well as anti-concentration results and tail estimates in a general context. This has been a very active trend of research in the last years, see for instance [17,18].

Declaration of competing interest

There are no competing interests.

Acknowledgements

We thank Oriol Serra for a preliminary reading of this manuscript, as well as suggestions and improvements in the presentation of the results. We also thank the anonymous referees of the paper, whose suggestions had improved the final presentation of the paper.

This research was conducted while Maximilian Wötzel was a member of the Barcelona Graduate School of Mathematics (BGSMath) as well as the Universitat Politècnica de Catalunya. Maximilian Wötzel acknowledges financial support from the Fondo Social Europeo and the Agencia Estatal de Investigación through the FPI grant number MDM-2014-0445-16-2 and the Spanish Ministry of Economy and Competitiveness, through the María de Maeztu Programme for Units of Excellence in R&D (MDM-2014-0445), through the project MTM2017-82166-P, as well as from the Dutch Science Council (NWO) through the grant number OCENW.M20.009.

Juanjo Rué acknowledges financial support from Spanish State Research Agency through projects MTM2017-82166-P, PID2020-113082GB-I00, the Severo Ochoa and María de Maeztu Program for Centers and Units of Excellence in R&D (CEX2020-001084-M), and the Marie Curie RISE research network ‘RandNet’ MSCA-RISE-2020-101007705.

References

- [1] K.F. Roth, On certain sets of integers, *J. Lond. Math. Soc.* 28 (1953) 104–109, <https://doi.org/10.1112/jlms/s1-28.1.104>.
- [2] E. Szemerédi, On sets of integers containing no k elements in arithmetic progression, *Acta Arith.* 27 (1975) 199–245, <https://doi.org/10.4064/aa-27-1-199-245>.
- [3] P. Frankl, R.L. Graham, V. Rödl, Quantitative theorems for regular systems of equations, *J. Comb. Theory, Ser. A* 47 (2) (1988) 246–261, [https://doi.org/10.1016/0097-3165\(88\)90020-9](https://doi.org/10.1016/0097-3165(88)90020-9).
- [4] Y. Kohayakawa, T. Łuczak, V. Rödl, Arithmetic progressions of length three in subsets of a random set, *Acta Arith.* 75 (2) (1996) 133–163, <https://doi.org/10.4064/aa-75-2-133-163>.
- [5] M. Schacht, Extremal results for random discrete structures, *Ann. Math. (2)* 184 (2) (2016) 333–365, <https://doi.org/10.4007/annals.2016.184.2.1>.
- [6] D. Conlon, W.T. Gowers, Combinatorial theorems in sparse random sets, *Ann. Math. (2)* 184 (2) (2016) 367–454, <https://doi.org/10.4007/annals.2016.184.2.2>.
- [7] D. Saxton, A. Thomason, Hypergraph containers, *Invent. Math.* 201 (3) (2015) 925–992, <https://doi.org/10.1007/s00222-014-0562-8>.
- [8] J. Balogh, R. Morris, W. Samotij, Independent sets in hypergraphs, *J. Am. Math. Soc.* 28 (3) (2015) 669–709, <https://doi.org/10.1090/S0894-0347-2014-00816-X>.
- [9] C. Spiegel, A note on sparse supersaturation and extremal results for linear homogeneous systems, *Electron. J. Comb.* 24 (3) (2017) 3.38, <https://doi.org/10.37236/6730>.
- [10] R. Hancock, K. Staden, A. Treglown, Independent sets in hypergraphs and Ramsey properties of graphs and the integers, *SIAM J. Discrete Math.* 33 (1) (2019) 153–188, <https://doi.org/10.1137/18M119642X>.
- [11] J. Rué, O. Serra, L. Vena, Counting configuration-free sets in groups, *Eur. J. Comb.* 66 (2017) 281–307, <https://doi.org/10.1016/j.ejc.2017.06.027>.
- [12] C. Spiegel, Additive structures and randomness in combinatorics, Ph.D. thesis, Universitat Politècnica de Catalunya, 2020, available at <http://hdl.handle.net/2117/328203>.
- [13] J. Rué, C. Spiegel, A. Zumalacárregui, Threshold functions and Poisson convergence for systems of equations in random sets, *Math. Z.* 288 (1–2) (2018) 333–360, <https://doi.org/10.1007/s00209-017-1891-2>.
- [14] A. Ruciński, When are small subgraphs of a random graph normally distributed?, *Probab. Theory Relat. Fields* 78 (1) (1988) 1–10, <https://doi.org/10.1007/BF00718031>.

- [15] Y. Barhoumi-Andréani, C. Koch, H. Liu, Bivariate fluctuations for the number of arithmetic progressions in random sets, *Electron. J. Probab.* 24 (2019) 145, <https://doi.org/10.1214/19-EJP391>.
- [16] I.Z. Ruzsa, Solving a linear equation in a set of integers. I, *Acta Arith.* 65 (3) (1993) 259–282, <https://doi.org/10.4064/aa-65-3-259-282>.
- [17] L. Warnke, Upper tails for arithmetic progressions in random subsets, *Isr. J. Math.* 221 (1) (2017) 317–365, <https://doi.org/10.1007/s11856-017-1546-3>.
- [18] R. Berkowitz, A. Sah, M. Sawhney, Number of arithmetic progressions in dense random subsets of $\mathbb{Z}/n\mathbb{Z}$, *Isr. J. Math.* 244 (2) (2021) 589–620, <https://doi.org/10.1007/s11856-021-2180-7>.